2
The Failure of Analysis and the Nature of Concepts

Michael Huemer

Introduction

Over the last century, many well-qualified philosophers spent many years attempting to analyze philosophically interesting concepts, such as KNOWLEDGE, FREE WILL, and CAUSATION. Yet no one succeeded in producing a single correct analysis. What went wrong? I ascribe the aspirations of conceptual analysis to a Lockean theory of concepts that ought to be rejected. I propose an alternative picture of concepts and properties that explains both (i) why linguistic intuitions about cases dominate the evaluation of conceptual analyses; and (ii) why most concepts are unanalyzable.

1 The failure of conceptual analysis: a story of twentieth century philosophy

The school of conceptual analysis has had its day. During the twentieth century, many philosophers viewed it as their central task to analyze words or concepts. Many spent decades discussing the analysis of such concepts as KNOWLEDGE, CAUSATION, and GOODNESS.[1] These were not incompetent people. As a rule, these philosophers were highly intelligent, were well-educated, and had about as strong a grasp on the concepts they studied as anyone. They had a great deal of time and attention for the task, and they worked in a community highly supportive of their project. If, therefore, the project of conceptual analysis should ever have been expected to bear fruit, it should have done so in the twentieth century.

It is therefore something of a shock to contemplate the meager fruits of the conceptual analysts' labors. To be sure, many subtle and valuable distinctions were drawn, and many analyses were shown to be *in*correct. But as to *correct* analyses, analyses immune from the ingenious-counter-example-generators in other philosophers' brains, the conceptual analysts have little to show.

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

No generally accepted analysis of any philosophically interesting term has yet been devised. Indeed, it is not out of line to maintain, as I believe, that no correct philosophical analysis has *ever* been produced in the history of philosophy.

There was a time when this shortage of results could have been chalked up to insufficient time or effort devoted to the task, or insufficient subtlety and cleverness on the part of the researchers. Perhaps at the beginning of the twentieth century, this would have been a reasonable line to take. But surely by now, any such apology has worn thin. We can never prove with certainty that the long-awaited analyses are not just around the corner. But at some point, one must, if one is reasonable, begin to question the basis for the search. Why did we think that our job was to analyze concepts? Why did we think that we could do it? More importantly, why has the endeavor met with so little success, and what can we learn about the nature of concepts from reflecting on the difficulty of analyzing them?

2 The roots of conceptual analysis

We all know the basic story. The logical positivists, inspired by Hume, assumed that only two kinds of knowledge were possible: analytic *a priori* knowledge, and synthetic empirical knowledge. Noticing that philosophers did not seem to be performing any experiments or making any specialized observations, they concluded that philosophy was not an empirical discipline; it must therefore be devoted to analytic knowledge, which, in the positivists' view, was entirely grounded in the meanings of words or concepts. Therefore, it must be the job of philosophers to analyze words or concepts.

This is not the place to criticize the positivist and empiricist dogmas.[2] What interests me here is another theory, which led to the expectation that the task of analyzing concepts should be relatively tractable. This is the classical theory of concepts, espoused by the likes of David Hume and John Locke.[3] These thinkers believed that concepts (or 'ideas') were introspectively observable mental items, which could be divided into two categories: *simple* concepts and *complex* concepts. Simple concepts were directly based on experience, either sensory or introspective. Locke and Hume viewed this process in an overly imagistic manner, treating simple concepts as copies of sensory images present in perception. Complex concepts, which comprised the large majority of concepts, were held to be constructed out of simple concepts, thus containing simple concepts as components. A concept's definition would reveal how the concept was constructed while simultaneously providing the criteria by which one applied the concept in particular cases.

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

Three tenets of this Lockean theory of concepts are of particular interest here:

1. Concepts are open to direct introspective examination.
2. Most concepts are composed of other concepts.
3. Definitions govern the application of concepts.

If these claims are true, we should expect conceptual analysis to be fruitful and valuable. If (2) is true, then most concepts will have analyses. If (1) is true, we should be able to discover these analyses, with a reasonable degree of reliability, by examining our concepts and seeing how they are composed. If (3) is true, these analyses will be useful in determining how to apply our concepts.

The history of twentieth-century philosophy poses a challenge to the Lockean theory of concepts, in at least two ways. First, the Lockean theory cannot explain the extreme difficulty of correctly analyzing concepts. The best-known illustration of this difficulty is the enormous literature on the analysis of KNOWLEDGE, inaugurated by Gettier's counter-examples to the justified-true-belief analysis. Fifty years and scores of counter-examples later, philosophers still have not agreed upon the correct analysis of the concept.[4] Much the same is true of every other philosophically interesting concept that philosophers have sought to analyze, including the concepts expressed by the words 'if', 'can', 'cause', 'good', and so on. Why have philosophers not been able to simply inspect their concepts and see what constituents they contain, in the same way, for example, that I can inspect my mental image of a horse and see that it contains parts representing legs, torso, head, and so on?

Second, the Lockean theory does not explain why philosophical practice, in the course of evaluating conceptual analyses, has been so strongly driven by *examples*. Before Gettier, philosophers were inclined to think that knowledge could be defined as justified, true belief. When confronted with Gettier's examples of justified but only accidentally true beliefs, we do not say, 'These cases satisfy the definition for KNOWLEDGE that we have accepted; therefore, they are cases of knowledge.' Rather, we see that the cases are not knowledge, and we deduce that the previously accepted definition is wrong. On the classical theory of concepts, we decide whether to apply a concept to an object by comparing the object's properties with the properties listed in the concept's definition. But philosophical practice suggests the reverse: we intuitively judge whether a concept applies to an object, independent of any definition, and we evaluate a definition by how well it fits with the correct usage of the concept.

The most plausible explanation for these two aspects of twentieth-century philosophical history begins with the admission that the Lockean theory of concepts is wrong. This should come to us as little surprise; very few theories

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

that held sway in the 1600s are still regarded as basically accurate today. We must look for an alternative view of concepts that explains both the difficulty of conceptual analysis and the role of intuitions about cases in evaluating proposed definitions. Once we have a more plausible theory of concepts, we may reconsider both the feasibility and the intellectual value of the project of conceptual analysis in the light of that theory.

3 Concepts, properties, and natures: a non-Lockean picture

3.1 Properties as regions

Here is an alternative theory. I do not know whether it is correct, but I shall hereafter assume the role of advocate to see how far the theory can be defended. The theory begins with a view about properties.

Every particular in the world – whether an individual object, state of affairs, event, or other particular – has a specific *nature*. This nature is a *comprehensive* and *fully determinate* property – that is, it includes everything (qualitative) about the way the object is. Every particular must have exactly one nature. These natures have no names in natural language; hence, I can refer to them only by such descriptions as 'the nature of this pen' and 'the nature of this pain sensation'. Natures are universals, in the sense that they can in principle be predicated of more than one thing; two objects that share a nature would be qualitatively identical. In addition to the natures instantiated by actual particulars, there are many possible natures that are uninstantiated – that is, there are ways that things could have been, such that no things actually are those ways.

In addition to natures, there are *abstract properties*. These properties are not comprehensive; they constitute only a part or aspect of how an object is. Most are less than fully determinate. For example, the abstract property *red* is only one aspect of a red object's qualitative nature. *Red* is less than fully determinate, since there is a range of shades that count as red – that is, *red* is a determinable of which various shades are determinates.

What is the relationship between natures and abstract properties? A traditional metaphysical view would take certain abstract properties as metaphysically fundamental. Specific natures would be understood as constructed out of the fundamental abstract properties. A nature might be described simply as a maximal conjunction of abstract properties – that is, when you conjoin enough abstract properties together, such that no further properties could be added without generating a contradiction, then you have a specific nature.

That is one possible view, but not the only view one could take. To explain an alternative, I begin with an analogy involving color, taking color as we think of it intuitively, color as it appears to us (thus, for the sake of exposition,

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

ignore reductionist and eliminativist theories of color). Every colored object has some specific, fully determinate shade of color. The specific shades can be arranged into the color cylinder, where hue varies as one moves around the circumference of the cylinder, saturation varies as one moves along a radius of the cylinder, and luminance varies as one moves up or down the height of the cylinder (see Figure 1). The arrangement of the points in this 'color space' corresponds to objective, internal similarity relations among the specific shades.

Now, what is the relationship between the specific shades and the hue, saturation, and luminance values? One answer is that each shade simply consists of a certain combination of hue, saturation, and luminance. On this view, hue, saturation, and lightness values are metaphysically fundamental, with specific shades being constructed out of and metaphysically dependent on those more abstract properties. Thus, a specific shade might be identified with the combination of a certain hue in the red range with a 60% luminance value and a 30% saturation value.

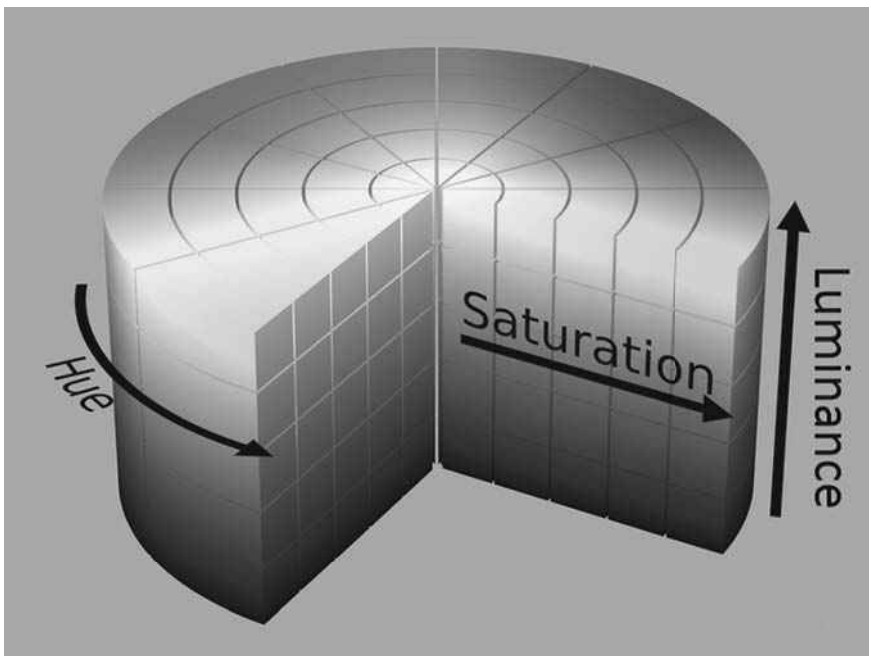

Figure 1 The color cylinder

Source: Creative Commons. 'The HSV color model mapped to a cylinder' by SharkD is licensed under CC BY-SA 3.0 (http://creativecommons.org/licenses/by-sa/3.0/deed.en).

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

But that picture seems backwards. Abstractions such as *30% saturation* are not metaphysically fundamental. It is the specific, determinate shades of color that are (more) fundamental. Abstract properties are just *regions* in the color space.[5] For instance, the abstract property *30% saturation* corresponds to a certain cylindrical surface within the color solid (with a radius 30% of the radius of the whole color solid). Likewise, *60% luminance* corresponds to a certain circular cross section taken 60% of the way up the height of the color cylinder. Each determinate hue corresponds to a vertical plane in the cylinder, extending from the center to the circumference. A determinable color such as *red* is a wedge-shaped region in the color cylinder (see Figure 2). Thus, rather than particular shades being formed by conjoining abstract properties, it is the abstract properties that are formed by grouping together specific shades.

One way of motivating this view of colors is to reflect on the different ways of classifying colors. Any specific shade can be conceptually described in terms of a triple of hue, saturation, and luminance values. But the same shade can also be described, instead, in terms of a triple of red, green, and blue components. Thus, for example, your computer monitor will display a certain paradigmatic bright orange, which can be described as (255, 128, 0) in the red–green–blue scheme, or (20, 240, 120) in the hue–saturation–luminance scheme.[6] Each scheme for classifying colors covers the same range of possible colors, and each scheme maps each color onto one of 16.7 million ordered triples. If we think of the abstract properties as the fundamental, objective components of a specific color shade, we can then ask whether a particular color shade is really constituted by its hue, saturation, and lightness properties, or by its red, green, and blue components. Which are the truly fundamental constituents of all colors? But clearly this question is senseless. Neither scheme for classifying colors is any more correct or metaphysically fundamental than the other. What is there,

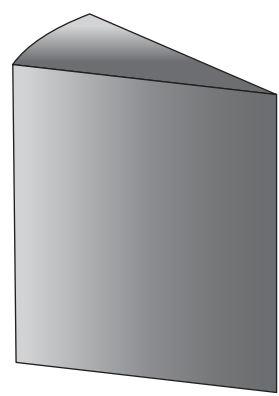

Figure 2 The color red as a wedge in the color cylinder

fundamentally, is just the various specific shades, with their internal relations to each other, and it is up to us to decide how to categorize them. The HSL and RGB conceptual schemes simply group shades together in different ways.

Now, perhaps what goes for colors goes for properties in general. Perhaps it is the determinate natures of things that are fundamental, and all the various abstract properties are constructed from these natures. The world does not come to us divided into abstract properties, such as *red*, *round*, *happy*, and so on. The world comes to us as a collection of objects each with its own precise, complete nature. It is up to us to *group* natures together into categories, thus delineating the abstract properties we choose to recognize. Abstract properties may exist independently of us (just as all manner of odd regions of space exist independently of us), but it is up to us which of the infinitely many properties will attain recognition in our conceptual schemes.

Just as shades of color may be regarded as points in a metaphorical 'color space', specific natures may be thought of as points in a metaphorical 'property space'. Natures have internal similarity relations to each other, thus giving rise to a notion of distance in the space. The more similar two natures are to each other, the 'closer' they are in the space. Thus, a particular pen's nature is closer to that of another pen than either is to, say, the nature of a particular dog. (In suggesting this metaphor, I do not mean to suggest that the mathematical properties of the space of possible natures mirror those of physical space. I leave open the question of the mathematical structure of the space of natures.) In terms of this metaphor, then, an abstract property is a *region* in the space of possible natures.

This theory is reminiscent of a certain version of nominalism, which identifies properties with classes of mutually resembling particulars.[7] The theory is not a form of nominalism, however, since it identifies abstract properties, not with collections of particulars, but with regions (or disjunctions) of mutually resembling *natures*, where each of these natures is a universal.[8] The view is thus not committed to denying the existence of universals, nor to denying the existence of abstract properties; it claims only that specific natures are more fundamental than abstract properties.

3.2 Conceptualization as boundary-drawing

Given the above view of properties, how should we understand *concepts*? Concepts are intentional mental states that represent abstract properties. Forming a concept should be understood as a matter of drawing a boundary around a region in the space of natures, grouping together all the natures in that region and distinguishing them from everything outside that region. Just as in the case of physical space, there are infinitely many regions in the space of natures; however, only a limited number of regions will be *recognized* in any human conceptual scheme. Most regions, that is, will fail to correspond

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

to any actual concept. Different conceptual schemes will draw boundaries in different places. These different ways of drawing boundaries are neither right nor wrong, though some are more useful than others.

When we categorize the objects we see around us – 'there is a leaf', 'there is a duck', and so on – we are constructing a much lower resolution picture of the world around us than the picture given to us by sensory perception. Perceptual experience makes very fine distinctions, giving us very detailed information about the world. Conceptual schemes simplify the world, treating objects that fall within certain relatively broad ranges as relevantly the same. Why are such low-resolution pictures of the world useful? One reason derives from the way in which they enable us to bring past experiences to bear on present circumstances. When I walk into a room, I encounter an object whose specific nature I have never before encountered. If I had no conceptual scheme at all, I would simply see this as a new phenomenon. It is because I have a pre-established category, TABLE, into which I can fit the new object that I have some expectations about this object, based upon my background knowledge and experiences involving tables. I expect the object to persist over time, to be capable of supporting smaller objects, and so on. Another benefit of conceptual schemes is the ability to communicate. We cannot have labels for every specific nature. If human beings are to communicate information to each other, we must have labels that correspond to certain broad regions in the space of natures. Conceptual schemes may also lighten the cognitive loads on our minds. Just as it is easier to store and manipulate a low-resolution image than to store and manipulate a very high-resolution image, it is easier to remember and reason about conceptualized information than to remember and reason about the sort of extremely complex and detailed information contained in sensory experiences.

3.3 Concepts as dispositions

What determines the boundaries of a concept?[9] Here is a naive view: the boundaries of a concept are determined by beliefs of the subject that identify those boundaries. For instance, perhaps the boundaries of the concept RED are determined by object-level beliefs to the effect that things with certain color shades are red. But this could not be correct, since one must first possess the concept RED before one can have any beliefs about what is red, and thus the concept must already have some independently set boundaries. (The nature of concepts precludes one's having a concept without boundaries.)

Alternately, perhaps the boundaries of the concept RED are determined by semantic beliefs, such as the belief that the concept RED applies to things falling within a certain range of hues. But this view also seems untenable. Beliefs are conceptual states; to believe that things of kind J stand in relation R to things of kind K, a subject must possess concepts for J, R, and K. So to believe

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

[the concept RED applies to things within color-range R], one must possess (i) the concept THE CONCEPT RED, (ii) the concept APPLIES TO, and (iii) the concept of color-range R. It is therefore not plausible that such a belief is needed to set the boundaries for the concept RED. Indeed, it is implausible in general that the boundaries for the concept RED are set by any conceptual mental state.

Perhaps our question is misguided. Perhaps – at least in the case of some concepts, the 'primitive' concepts – we should simply accept a concept's boundaries as a fundamental, inexplicable fact. But while it may not be possible to provide a complete and noncircular account of the boundaries of a concept, some interesting and informative remarks can be made. When one possesses a conceptual category, one is disposed to view the things falling under that category in a certain way – to see them as alike in a certain way, and as different from the things outside the category. The contours of the category are determined by these dispositions: how far a concept extends is a matter of what range of natures the concept-wielder is disposed to classify together, to see as 'alike'.[10] This holds true at least for primitive concepts, for which the subject lacks definitions. This dispositional character of concepts enables them to determinately apply or fail to apply even to cases that the subject has never considered, since a subject may still possess determinate dispositions in regard to such cases. However, a concept may also have indeterminate boundaries, since a subject may also fail to possess determinate dispositions in regard to certain cases – that is, there may be some things such that a subject is neither disposed to apply the concept, nor disposed to reject the application of the concept to those things (this is true of borderline cases for vague predicates).

3.4 The sources of conceptual dispositions

Where do our conceptual dispositions come from? What makes us classify some things together, and not others? Objective similarity plays an important role; we dislike categories that exclude some objects that are more similar to some items in the category than the items in the category are to each other.[11] Most people strongly resist forming the concept GRUE for precisely this reason. GRUE includes blue objects observed before the year 2100, and green objects not observed before 2100, but it excludes blue objects not observed before the year 2100.[12] Most people find this capricious, because blue objects not observed before the year 2100 are *more similar* to blue objects observed before the year 2100 than either is to *green* objects not observed before 2100. The concept GRUE thus defies our similarity-based classification practices.

But there must be other influences on our categorization practices. The similarity principle provides only a very weak constraint (and one which we may be willing to violate in some cases). One such influence is provided by one's language. The individual is influenced by the other members of his speech community, to form concepts with boundaries corresponding to the usage of

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

particular lexemes in that community. When one hears the word 'know', for example, applied to a given case, one is influenced toward applying the word in other cases similar to that case. One's concept of knowledge is largely a resultant of all these influences. These influences are very difficult to overcome, so that individuals who attempt to create stipulative usages of existing words almost always continue to be influenced by the standard usage of those words in their speech communities.

Another common tendency, in normal circumstances, is the tendency to form concepts that are useful and avoid concepts that are useless. Thus, among other things, we want our conceptual schemes to identify similarities and distinctions that are relevant to our interests. This desideratum takes us beyond the classification of objects according to their intrinsic similarities. For instance, consider the deaths of Smith, Jones, and Brown: Smith was killed by an avalanche, which was intentionally started by another human being. Jones was killed by an avalanche *not* started by another human being. Brown was murdered by poison. If the coroner knows the facts of these cases, he will classify Smith's and Brown's deaths as 'homicides', whereas Jones' death was a 'death by natural causes' – despite the fact that Smith's and Jones' deaths (including the events proximately causing each) were more physically similar to each other than either was to Brown's death. The reason for this classification scheme is practical: we care much more about whether a death was caused by human action than we do about whether it was caused by an avalanche.

In line with the desire for useful concepts, we typically seek to form concepts that are informative, in the sense that a person who learns that concept C applies to X thereby learns a good deal. The sort of informativeness that is practically relevant, and that will make a concept useful, is probabilistic rather than deductive. That is, a person who learns that C applies to X should 'learn a lot' in the sense of being induced to make a significant revision to his subjective probability distribution, rather than in the sense of being able to deduce many theorems. Informativeness is thus relative to a prior probability distribution. Keeping in mind the mutual influence of members of a single speech community on each other's concept formation, we should anticipate, roughly, that people will tend to form concepts that would be informative for most members of their speech community.

At the same time, however, we want our concepts to have reasonably wide application, and this desideratum stands in some tension with the desideratum of informativeness. The most informative concept would be one whose application to an object locates that object at a specific point in the space of possible natures. Such a concept, however, would be of little use, since in all probability, no more than one object would ever fall under it. We tend, therefore, to form concepts that apply to many of the objects that we observe and care about.

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

Most of these rules are vague, and none is inviolable. Nor have I tried to give a complete list of influences on concept formation. These are simply some of the more prominent influences on our categorization practices.

I want to draw attention to one special case that I think is particularly common and important. Suppose we find a class of actual objects that form a cluster in the space of possible natures – that is, each of these objects is relatively similar to some other objects in the class, and relatively dissimilar from any actual objects not in the class. Then we are likely to group those objects together under a single category. Our observations of objects in the world thus influence where we draw conceptual boundaries (see Figures 3, 4).

An interesting example of this tendency is provided by the 2006 reclassification of Pluto by the International Astronomical Union. It was once said that there were nine planets in the solar system, of which Pluto was the outermost and smallest. Then astronomers began to identify other objects that are similar to Pluto and relatively dissimilar from the eight remaining planets (in the same way that Pluto is dissimilar from the other eight planets). As a result, if Pluto continued to be classified as a planet, we would 'have to' add fifty more objects to the list of planets, all of them much smaller than the first eight planets, and all but one of them located in the Kuiper belt at the edge of the solar system. This discovery made it more natural to place Pluto and the other Pluto-like objects into their own category, that of 'planetoids', thus decreasing the number of 'planets' to eight (see Figure 5).[13] Had Pluto been the only object of its kind, it would probably have retained the title of 'planet'. It was Pluto's bad

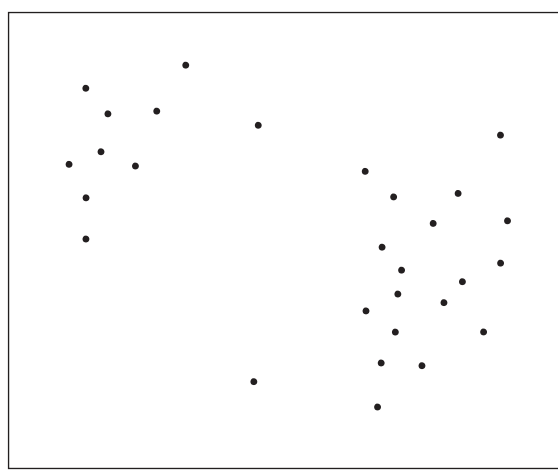

Figure 3 Clusters of objects in the space of possible natures. Each point in the space represents a nature that something could have. Dots represent natures of objects actually found in the world

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

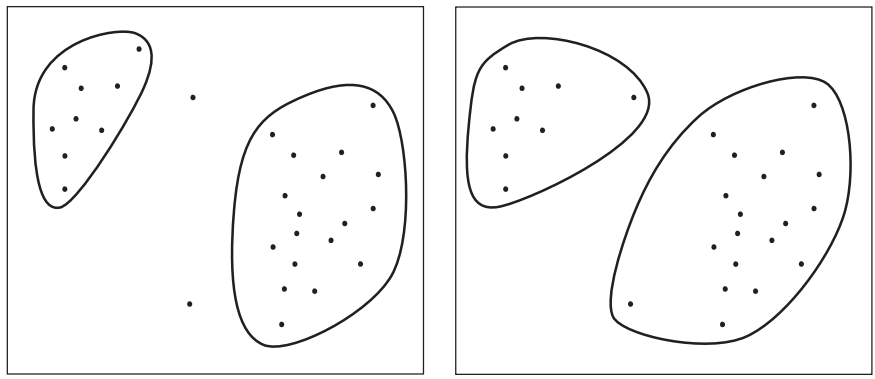

Figure 4 Two natural ways of grouping clustered objects into conceptual categories

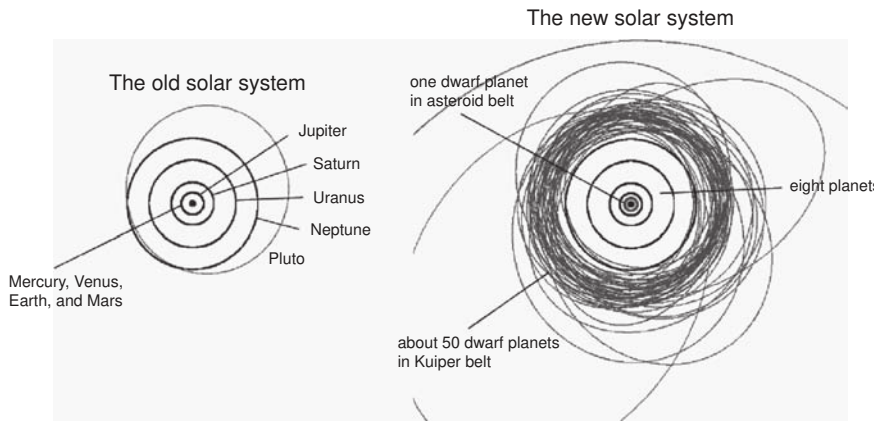

Figure 5 Evolving knowledge of the solar system led to the reclassification of Pluto from a 'planet' to a 'dwarf planet' or 'planetoid'

luck to fall in with a large cluster of objects, all resembling each other much more than any of them resembled any of the original eight planets.[14]

4 Implications for analysis

A successful conceptual analysis of the kind philosophers have sought for the last century must at least provide logically necessary and sufficient conditions for a thing to fall under the concept that is being analyzed. Thus, suppose we wish to analyze KNOWLEDGE. A successful analysis of KNOWLEDGE would

provide a set of conditions that pick out all and only the mental states to which KNOWLEDGE applies, in every coherently conceivable circumstance. For such an analysis to be possible, we must possess *other* conceptual categories that are related in a very specific way to the category of KNOWLEDGE – roughly speaking, it must be possible to construct the region in the space of natures corresponding to KNOWLEDGE through set theoretic operations, such as union, intersection, and subtraction, on the regions corresponding to some other concepts that we possess. For example, suppose that the region in the property space corresponding to KNOWLEDGE happens to be identical to the intersection of the regions corresponding to BELIEF, TRUTH, and JUSTIFICATION (see Figure 6). In that case, and only in that case, it will be possible to analyze KNOWLEDGE as JUSTIFIED TRUE BELIEF.

The first thing to notice is that the conditions under which this sort of analysis is possible are extremely restricted. Given that we possess only finitely many other concepts, there are only finitely many regions that can be constructed through set theoretic operations on those concepts. But there are uncountably many regions in the space of natures. So, for any randomly chosen region, the prior probability that that region can be precisely picked out using concepts that we already possess is zero (or infinitesimal). In other words, of all the regions that our category of KNOWLEDGE *could* correspond to, only a measure-zero subset of regions would be such that the concept would be analyzable in terms of other concepts that we possess. There is therefore no obvious reason to presume that the concept KNOWLEDGE is analyzable.

Of course, some concepts are analyzable in terms of other concepts. I am aware of exactly one area in which correct conceptual analyses are commonly found – that is mathematics. It appears that in mathematics, it is actually common for a concept to be constructed, or at least constructible, out of other

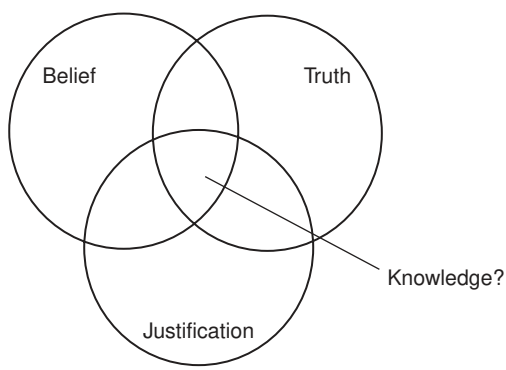

Figure 6 The relationship between BELIEF, JUSTIFICATION, and TRUTH in which KNOWLEDGE could be defined as JUSTIFIED, TRUE BELIEF

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

concepts. Mathematics, however, is a very peculiar area of human endeavor, and the relative difficulty that the majority of people have with understanding mathematical ideas and mathematical reasoning should tip us off to the possibility that mathematical thinking is quite different from ordinary thinking, and thus that we should be very wary of taking mathematics as a model for the rest of human cognition.

If we consider the sources of conceptual dispositions discussed above (section 3.4), we find very little comfort for the aspiration of providing conceptual analyses. If the boundaries of most ordinary concepts are determined by such factors as our interests, the desire for informativeness, the desire for wide applicability, and the desire to include clusters of observed objects, then ordinary concepts are very unlike the concepts of mathematics. Most ordinary concepts may be expected to correspond to regions with very complex and idiosyncratic shapes, and there is no obvious reason to expect them to be definable in terms of other concepts. In addition, if our concept-formation is strongly influenced by our language, it would be unsurprising if people from different linguistic communities used at least slightly different conceptual schemes.[15]

There is also no particular reason to expect conceptual analysis to be *interesting*. Even if we managed some day to construct some complicated formulation that succeeded in picking out all and only the cases of knowledge in every possible world, it would probably be a tediously baroque and gerrymandered formula, and it would probably do nothing to help anyone understand KNOWLEDGE. One does not form the concept KNOWLEDGE by reading a definition. One forms the concept by acquiring certain classificatory dispositions, largely through observation and imitation of other people's usage of the word 'know'. Any definition able to avoid all the counterexamples philosophers have devised to earlier definitions of KNOWLEDGE would almost surely contain so many complicated, abstract clauses that most readers would find it far more difficult to comprehend the definition than to comprehend the concept of knowledge itself. That definition would likely serve more to confuse than to enlighten readers.

Nor have we much ground for confidence in the introspective accessibility of conceptual boundaries. Those boundaries, I have suggested, are determined by our classificatory dispositions, and there is no obvious reason to expect those dispositions to be directly introspectible as such. *Occurrent* mental states are usually directly introspectible, but mere dispositions are seldom accessible *as such* (that is, while they remain mere dispositions). The easiest way to access one's dispositions is to *activate* them. For instance, one becomes aware of what dispositional beliefs one had when those beliefs become occurrent. One discovers whether one is afraid of heights by getting into a high place, or at least imagining being in a high place. Similarly, the way we observe the classificatory dispositions that constitute a concept is by activating them: we

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

encounter, or imagine, a case, and then either are inclined to apply the concept or are inclined not to apply it. It is very difficult, perhaps impossible, for a person to identify all the relevant dispositions purely through abstract reflection on a concept; it is much easier to activate the dispositions by reflecting on particular cases, and then observe how one is inclined to classify the cases.

The picture of concepts and properties I have sketched thus accounts for both the role of cases in evaluating analyses and the general difficulty of constructing correct analyses.

5 Questions about natures

5.1 Analyticity

The present view of concepts and properties raises many questions, of which only a few can be discussed here. To begin with, among the critics of conceptual analysis are some who reject the notion of analytic truths altogether.[16] What implications does my view of concepts and properties have for the issue of analyticity?

If we want to retain the notion of analyticity, or some variation on it, we should not characterize the notion in terms of *definitions*. For example, we should not say that a sentence is analytic just in case the replacement of a term in the sentence with its definition results in a logical truth.[17] What we can say, however, is that a concept will often pick out a region in the space of natures which entirely encompasses the region picked out by another concept. For instance, the region picked out by RED is a proper part of the region picked out by COLORED (see Figure 7). This observation brings us at least within the neighborhood of the idea that 'All red things are colored' is analytic. More generally, analytic statements will be guaranteed true by the relationships among the regions of the property space picked out by the words in the statement, rather than their truth's depending on the locations of actual particulars in the space.[18]

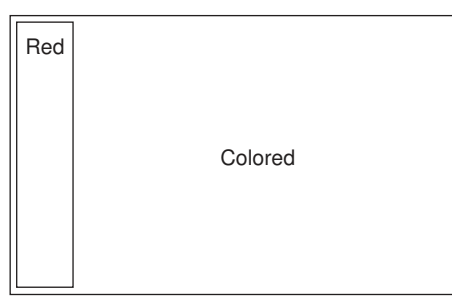

Figure 7 The red region as a proper part of the colored region

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

The foregoing may not quite capture what we intend by the term 'analytic'. Perhaps analyticity requires something more, something in the neighborhood of a sentence's providing a test of linguistic competence. Some say that analytic sentences are those for which one's linguistic competence, or one's understanding of the sentence, confers knowledge of or justification for accepting its truth.[19] All of this may be said consistently with the view of concepts and properties that I have proposed.

I shall not, however, attempt here to precisely analyze analyticity; any effort to do so would likely encounter the same sort of problems that afflict analysis in general. But how, one might wonder, can a person understand the concept of analyticity without an explicit definition? It is one thing to suppose that we may understand ordinary language terms without explicit definitions, but 'analytic' is clearly a philosophers' technical term. Aren't technical terms, at least, introduced through explicit definitions?

Consider how Kant initially introduced the term 'analytic'. Kant said that an analytic judgment is one in which 'the predicate B belongs to the subject A, as something which is (covertly) contained in this concept A.'[20] In the standard modern illustration of the idea, 'All bachelors are unmarried' is said to be analytic because the concept BACHELOR contains the concept UNMARRIED.

But this characterization clearly does *not* capture the meaning of 'analytic'. For on the above characterization, 'No bachelors are unmarried' would also count as analytic, since the concept of the subject (namely, BACHELOR) again contains the concept of the predicate (namely, UNMARRIED). Furthermore, 'No bachelors are married' would *not* count as analytic, since the concept of the subject (BACHELOR) does not contain the concept of the predicate (MARRIED).[21] And these are consequences that any philosopher competent in the use of the term 'analytic' would recognize as problematic. Thus, even though it was Kant who introduced 'analytic' as a technical term, we can still say that his definition of the term was wrong. This may seem odd. In defining a new technical term, doesn't one simply *stipulate* that the term means whatever one's definition indicates? How could Kant have gotten the definition *wrong*?

The answer is that Kant had a certain concept in mind, which he had formed by noticing a cluster of similar judgments (the judgment that all bachelors are unmarried, that everything is identical with itself, and so on), and he then *attempted* to describe the contours of that concept by giving the verbal formulation quoted two paragraphs above. But this formulation simply failed to capture the concept that Kant himself had in mind – just as twentieth-century philosophers who attempted to analyze KNOWLEDGE failed to correctly characterize their own concepts of knowledge. How can we tell that Kant did not really mean by 'analytic' what he said he meant? In this case, we can gather what Kant meant mainly from the role that analyticity played in his philosophy. For instance, Kant was concerned to solve the puzzle of how there can

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

be synthetic a priori knowledge. But if 'analytic' truly meant what Kant said it meant (with 'synthetic' understood as denoting the non-analytic), then the knowledge that no bachelors are married would be an instance of synthetic a priori knowledge. But we can be fairly confident that no philosopher would have made such heavy weather over how one can know such things as that no bachelors are married, *as opposed to* such things as that all triangles have three sides, as Kant did make over how we can have synthetic a priori knowledge, as opposed to analytic a priori knowledge.

One lesson is that to fix the meaning of a technical term, it is not necessary to correctly define that term. It may suffice to offer some examples and to make a series of statements and arguments using the technical term, where those statements and arguments make sense or are plausible only when the term is understood in the sense intended. Readers may then learn the use of the technical term in much the same way that they learn the use of terms in ordinary language. Of course, it might have been preferable had Kant given a correct definition of ANALYTIC (assuming a correct and comprehensible definition exists). But his failure to do so did not prevent his use of 'analytic' from being meaningful, nor did it prevent readers from attaining a basic grasp of the concept he had in mind.

5.2 Conceptual analysis and reduction

Some philosophers argue that conceptual analysis plays a key role in a certain metaphysical project, the project of providing reductions of higher-level phenomena to more fundamental phenomena. For instance, it is said that one can conclude that water is H_2O only after one has a conceptual analysis. One analyzes WATER along the lines of 'water is the watery stuff of our acquaintance', where 'watery' stands for a cluster of familiar properties of water (it is clear, odorless, and tasteless; it fills the rivers, lakes, and oceans on Earth; it falls from the sky as rain; and so on). Scientific investigation then reveals that H_2O in fact satisfies that definition of WATER, at which point we are in a position to conclude that water is in fact H_2O.[22]

On my view of concepts and properties, what becomes of this model of scientific reduction? Certainly the model must be revised. We cannot analyze WATER, so no such analysis can be required for discovering that water is H_2O. But we have not rejected the notion of analyticity altogether. It therefore may still be analytic that 'All watery stuff of our acquaintance is water', even if 'water' is not *defined* as 'the watery stuff of our acquaintance'. It is worth noting that among the objections to Jackson's quasi-analysis of water are counterexamples in which water would have had different observable properties from those it actually has.[23] These examples do not undermine the claim that, analytically, *if there is* some substance with which we are acquainted that has those properties, then that substance is water.

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

My view of concepts does not *commit* us to accepting this sort of model of reduction. But as far as I can see, it leaves open the possibility of a model of reduction in the spirit of the model proposed by Jackson, albeit with some revision.

5.3 The dimensions of the property space

If we understand abstract properties as regions in a space of possible natures, how are we to understand *dimensions*? Take for example the color space, which can be characterized in terms of dimensions of hue, saturation, and brightness. How should one understand the dimension of hue? Is this *dimension* also a region in the color space? If so, what region is it? We might say that it is the entire space, since every point in the space has some hue. But the same logic would lead us to say that saturation and brightness each correspond to that same region, and thus that hue, saturation, and brightness are identical.

Here is an alternative view. Just as we group together points in the color space to form particular hues, such as red, we group together particular hues to form the dimension of hue. The dimension thus is not simply a collection of points in the space, but a collection of collections of points. Hue differs from saturation in that the two dimensions group together different groups. As long as we accept that sets with the same ur-elements may nevertheless be distinct, we should have no difficulty in seeing hue, saturation, and brightness as distinct dimensions, despite the fact that exactly the same things have hues as have saturations and brightnesses.

Just as there are different ways in which a conceptual scheme may divide the possibility space into properties, there are also different ways in which a conceptual scheme may characterize the dimensions of the space. There is not a single correct set of dimensions for the space, any more than there is a single correct set of categories. To illustrate, consider again the color space, which may be conceptualized as a cylinder with hue, saturation, and brightness dimensions. This same space may also be conceptualized, equally correctly, as a cube with red, green, and blue components as dimensions.[24]

Even ordinary physical space admits of different ways of characterizing its dimensions. One may identify locations in space using either a Cartesian or a polar coordinate system, and in either case, one must stipulatively designate a special location (the 'origin') and two special directions (the '0 degree' directions or the 'x and y axes'). No point or direction in space is physically special, so these choices are epistemically arbitrary conventions, though some choices may be more practically useful than others.

5.4 Similarity and other relations

If the conception of properties sketched above is correct, it seems that something like it ought to apply equally well to *relations*. If an abstract property is a region in the space of possible natures, what is an abstract relation?

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

There are at least two interesting ways of treating relations that fall in line with the above treatment of properties. The first approach would be to treat a relation between two things as a property of a larger system – a relation between x and y is a property of a system (or an aggregate) containing both x and y. For instance, among the ways that the aggregate containing the cat and the mat might be is the arrangement in which the cat is on the mat. An abstract relation such as that expressed by 'x is on y' can be thought of as a certain grouping of points in a certain possibility space, a space whose points are all the comprehensive, fully determinate ways that a two-member system might be.

The other approach to relations is to treat relations as a separate but parallel phenomenon to properties. Just as we have recognized comprehensive, fully determinate *properties*, which I have called 'natures', we can also recognize comprehensive, fully determinate *relations*, which we might call 'relational natures'. Every pair of objects has exactly one relational nature – this is the nature that encompasses the complete, precise way that those objects are related to each other. These relational natures are also universals, in the sense that more than one pair of objects could in principle share a relational nature (in which case the members of each pair would be related to each other in exactly the same way, in every respect). We might then suppose that an abstract relation is a region in the space of possible relational natures, and that conceptualizing an abstract relation is a matter of drawing a boundary in that space.

We might wonder whether special difficulties are created by the relation of *similarity*, given the role that similarity plays in the rest of the theory. We have said that points in the possibility space bear objective, internal similarity relations to each other, with some pairs of points more similar than others. These similarity relations help to explain how we group points together when we form conceptual categories. But now suppose we recognize a space of relational natures, with abstract relations conceived as regions in the space. A given degree of similarity is itself an abstract relation. Can degrees of similarity therefore also be thought of as regions in this same possibility space? Presumably not. Degrees of similarity are the *distances* in the space; they cannot simultaneously be regions in the space.

There are two things that might be said about this. On the one hand, we might simply recognize similarity as an exception to the rule that properties and relations are regions in a possibility space. Similarity might be taken as fundamental, not subject to further explanation, with (most) *other* properties and relations understood as collections of mutually resembling natures.[25]

Alternately, we might view degrees of similarity as regions in *another* possibility space. Ordinary, first-order properties are regions in a space of ordinary, first-order natures, that is, the sort of natures that may belong to particulars. Similarity, one might say, is not a relation between particulars but a relation between natures; thus, it is a second-order universal. Second-order universals

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

might be understood as regions in a 'second order space', a space whose points are the comprehensive, fully determinate natures *of* first-order natures. A particular degree of similarity would thus be a particular region in the second-order space.

The latter way of treating similarity may be more in the style of the general approach to properties and relations that I have advanced. However, I find it less satisfactory overall, for two reasons: to begin with, I find the distinction between a first-order nature and the nature of that nature somewhat obscure. Second, the approach seems poised for an infinite regress, which I am not sure is not vicious. If the degree of similarity between two natures is to be understood in terms of a region in a second-order space, won't the degree of similarity between two second-order natures have to be understood in terms of a region in a third-order space (a space of natures of natures of natures)? For these reasons, it seems preferable simply to treat the degree of similarity of two natures as a fundamental fact, subject to no further explanation.

5.5 The individuation of particulars

In the preceding discussion, I have treated particulars, comprehensive natures, and the degrees of similarity between natures as given. But the identification of what particulars there are is as much a matter of human boundary-drawing as is the identification of what properties there are. The world of particulars does not come pre-divided for us, any more than the world of universals comes pre-divided. Can this observation be incorporated into our theory of properties?

An analogous view may be taken of the identification of particulars to the view we have taken of the formation of conceptual categories. In the case of the physical world, the analog of the space of natures is physical spacetime. Just as we form concepts by drawing boundaries in the relevant possibility space, grouping together the specific natures within those boundaries, so we identify particular physical things by drawing boundaries in spacetime and grouping together the material within those boundaries.

The two cases of course are not analogous in every respect. One difference is that our abstract concepts refer to regions in the possibility space, whereas our ideas of particular things typically refer not to regions of spacetime but to what occupies those regions.

The individuation of particulars, however, is not our primary concern. The question of relevance here is only whether our view of properties is undermined by its apparent need to rely upon a pre-established scheme for dividing the world into individual objects, events, and the like. On reflection, this does not appear to be a problem, since however we divide the world into particulars, those particulars will then have specific natures which can be grouped together to form abstract concepts. There is no need to assume a uniquely correct scheme for identifying particulars. Our scheme for identifying particulars will affect

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

what natures we see in the world – to take a trivial example, if we adopted a scheme in which only elementary particles were recognized, then all the recognized particulars would be very small, and thus would be located in a certain very narrow region of the space of possible natures. The particulars that we recognize will profoundly affect how we are inclined to group together natures into conceptual categories, since we form concepts partly on the basis of where we observe particulars to be located in the space of possible natures.

5.6 The awareness of specific natures

The notion of forming abstract concepts by grouping together determinate natures seems to suggest that we are pre-conceptually aware of the specific natures that things have or could have had. But this suggestion is problematic, since natures, as I have described them, are fully comprehensive and fully determinate properties; thus, to be aware of any object's nature, we would have to be aware of *exactly* how that object is, in every (qualitative) way. Needless to say, there are few if any things for which we have that kind of perfect awareness. One might be tempted to argue that we need only the awareness of a *possible* nature, since abstract concepts pick out regions in the space of possible natures, and thus that we need not know, of any specific nature, whether it is actually instantiated. But this suggestion makes little headway, since it is not particularly plausible that we are aware of specific natures even as mere possibilities. An ordinary human being probably cannot even imagine any of the complete and fully determine ways that ordinary objects might be.

Consider an illustrative analogy. A professor teaching a large lecture course asks one of his teaching assistants to divide the class into four-person groups for purposes of a group project that the professor plans to assign. But suppose that the teaching assistant lacks access to the course roster and knows at most only a few of the students in the class. In this case, it seems that the teaching assistant will not be able to perform the requested task. Without some way of identifying the individual students in the class, the TA will not be able to group students together. Similarly, one might think, without some way of identifying (a non-trivial range of) specific natures, we have no way of grouping these natures together.

Two observations help us to see our way around this puzzle. The first is that awareness of x does not typically require complete or maximally precise awareness of x. Suppose, for example, that I look at a penguin under normal conditions. I see the portion of the penguin's surface that faces me. I do not see (nor am I visually aware of) the back of the penguin or the inside of the penguin. And the penguin's back and inside make up most of the penguin. Nevertheless, it is correct to say that I see (and thus am visually aware of) the penguin. This imperfect awareness is enough for me to refer to the penguin, to group the penguin together with other penguins, and so on. In the same sense, it is

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

possible to be aware of an object's nature despite being unaware of some parts or aspects of that nature.

The second observation is that the way in which we group natures together when we form a concept is not like the way one might group together students in a class. In the teaching-assistant example, the professor presumably wants the TA to divide the class into groups by producing an explicit list of groups, with the names of each member of each group. This requires a list of who is in the class. If concepts consisted of explicit lists of everything included within them, then it would be for all practical purposes impossible to form a concept. Concepts, however, do not group things together in that way. Concepts group things together by means of dispositions. That is, the sense in which two natures are 'grouped together' under a concept is that the subject is *disposed* to see things with those natures as relevantly alike, as falling into the same category, when the subject becomes aware (in the limited way in which we often do so) of objects having those natures. Concept formation thus does not require the formation of long explicit lists, nor does it require knowledge of any complete nature.

6 Methodological conclusions: clarification or analysis

I am not confident that the theory of properties and concepts that I have outlined above is entirely correct. Probably, there are some errors in it. Nevertheless, I have a fair degree of confidence that the anti-conceptual-analysis lesson is essentially correct. Even if my theory contains errors, I think that enough of it is close enough to the truth to render the endeavor to analyze concepts unpromising and uninteresting. Perhaps specific, determinate natures are not really metaphysically prior to abstract properties. Even so, I think it would remain plausible that conceptual boundaries are set by classificatory dispositions, and highly unlikely, in light of everything we know, that most concepts are constructed out of other concepts. This is enough to undermine our confidence in the feasibility of conceptual analysis, as well as the assumption that analyses are important for understanding our concepts.

If the purpose of conceptual analysis was to clarify our concepts, this aim need not be abandoned; it merely need be pursued in a different manner. One can say many things to clarify a concept, short of offering a philosophically satisfying definition. The first and most obvious thing one can do is to offer examples in which the concept is appropriately deployed. These examples should make some effort at indicating the breadth of the concept. For instance, to explain the category BIRD, one would mention several very different sorts of birds, such as sparrows, eagles, ostriches, penguins, and chickens.

Second, one can distinguish the category in question from other categories that are related to but distinct from it. For instance, suppose one wants to

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

clarify the notion of an *excuse*. One would want to distinguish excuses from justifications and simple denials. One could give examples in which a person offers an excuse for some behavior but neither denies nor justifies the behavior, other examples in which a person justifies rather than excuses some behavior, and still others in which a person denies rather than justifies or excuses the behavior.

Third, one can give verbal formulations that *approximate* the meaning of a target term. Thus, we might say that roughly speaking, an excuse is an attempt to explain why one should not be blamed for some action, without challenging either that one performed the action or that the action was unjustified. This is something *close* to an analysis, but there is no suggestion that this formulation would stand up to all logically possible counter-examples. Despite its merely approximate character, this explanation of EXCUSES may help clarify the concept for those who initially fail to grasp it.

Fourth, one can often give conditions that are either necessary or sufficient for a concept's applicability even when one cannot give a single set of conditions that are necessary and sufficient. Thus, some clarification of KNOWLEDGE is provided by noting that a person knows that P *only if* the person at least believes that P and P is true. Likewise, one can note that *if* a person has a true belief that P and there are no defeaters for that belief, *then* the person knows that P.[26]

Fifth, one can clarify a concept to some degree by discussing such matters as the role that the concept plays in human life, why the category established by that concept is important, and what are the further implications of something's falling under that concept.[27] For instance, perhaps knowledge functions as a kind of doxastic ideal: knowledge is what one should be aiming for when forming beliefs; a belief that is fully successful is knowledge.

Finally, one can clarify a concept (or the aspect of the world to which it refers) by taxonomizing its referents, by discussing the logical features of the concept, and by examining other general features of the phenomenon to which the concept refers. For instance, one may clarify KNOWLEDGE (or at least knowledge) by discussing such matters as the types of knowledge that exist, how knowledge comes about, whether knowledge is closed under entailment, and so on.

Some of this discussion would be a priori (perhaps it is knowable a priori whether knowledge is closed under entailment) while some would be empirical (we know empirically that testimony is a source of knowledge). Perhaps some philosophers would wish to exclude empirical observations from any project to be construed as conceptual clarification. One reason why this would be a mistake is that there is no sharp dividing line between clarifying concepts and clarifying reality, and empirical discoveries often have a way of doing both simultaneously. Prior to the discovery of the Theory of Evolution, for example, it would not have occurred to anyone that the membership of such

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

biological categories as BIRD, FISH, and MAMMAL should be determined by common ancestry. We could describe this as a conceptual fact – it is, after all, an observation about the *classification* of certain phenomena, about how these things are to be grouped together. But it is a fact in whose recognition empirical science had a large role to play.

Moreover, it is difficult to see the urgency of separating conceptual clarification from illumination of the world, even if we could draw this distinction with confidence in all cases. Concepts are tools for understanding reality, and the purpose of clarifying concepts, presumably, is to improve our understanding of reality. It is unclear what cognitive goal requires us to focus on a priori conceptual clarification while carefully avoiding taking into account any empirical facts.

In seeking to clarify our understanding of the world, we philosophers have construed our task too narrowly. We have assumed that the task consists chiefly in saying precisely what a concept means. But 'saying what something means' simply involves us in deploying *other* concepts, the concepts expressed by the words in the stated explanation. So the endeavor is in fact one of *translating* a concept into *other* concepts that supposedly constitute it. But conceptual schemes are not structured in this way; concepts are typically not composed of other concepts nor understood through operations on other concepts. This is the lesson of the failure of analysis. The best way of improving our understanding is, therefore, simply to clarify the nature of the properties and relationships that our concepts pick out.[28]

Notes

1. I denote concepts with all-upper-case letters.
2. For criticisms of empiricism, see BonJour (1998); Bealer (1992); Huemer (2005, chapter 5).
3. Hume (1975, section II); Locke (1975, book II).
4. Gettier (1963). For a review of the first twenty years of the literature, see Shope (1983). Williamson (1995, pp. 541–543) takes the history to evidence the unanalyzability of KNOWLEDGE.
5. Gardenfors (2000, esp. chapter 3) defends this idea, both for the case of color and for the case of properties in general.
6. These are the two schemes used for identifying colors in many computer programming contexts. 0 is the minimum value, and 255 is the maximum value on each dimension, whether one uses the HSL or RGB scheme.
7. Price (1969); Rodriguez-Pereyra (2002).
8. The abstract property should be thought of as something like a *disjunction* of specific natures, since an object instantiates a given abstract property precisely when it has any of the specific natures in the relevant region.
9. I use ' determine' here to express a constitutive or *in virtue of* relation, rather than either a causal or an epistemic relation.
10. For a similarly dispositional view of concepts, see Wilson (1982, pp. 567–569).

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly

11. A more general principle is that we prefer concepts that correspond to convex regions in the property space, regions such that for any two points in the region, every point located between those two points is also in the region. Gardenfors (2000, pp. 70–77) argues that only such regions count as natural properties, as opposed to gerrymandered properties. Cf. Oddie (2005, pp. 152–158).
12. Goodman (1955), chapter III.
13. Figure 5 and the account of the classification of Pluto come from Caltech astronomer Mike Brown, <http://web.gps.caltech.edu/~mbrown/dwarfplanets/>.
14. For other examples of the historical contingency of concept extensions, see Wilson (1982, pp. 549–551, 572–574).
15. See, for example, Weinberg, Nichols, and Stich's (2001, pp. 439–448) evidence suggesting that Asians may have a slightly different concept of 'knowledge' from Westerners. See Goddard (2001) for a review of the evidence concerning which word meanings, if any, are universal among human cultures.
16. Quine (1951).
17. This conception of analyticity derives from Frege (1980, section 3).
18. Gardenfors (2000, p. 166) advances this interpretation of analyticity. How to distinguish analyticity from necessity, however, remains unclear.
19. Boghossian (1996); Rey (2004).
20. Kant (1965, B10).
21. Unless, of course, we say that BACHELOR contains UNMARRIED, which itself contains MARRIED, and therefore BACHELOR contains MARRIED. But this would lead to our counting 'All bachelors are married' as analytic on Kant's definition.
22. Jackson (1994).
23. Laurence and Margolis (2003).
24. See Tkalcic and Tasic (2003) for a review of several different versions of the color space.
25. Presumably the relation of identity would be another exception.
26. See Klein (1971) on the relevant notion of defeaters, though Klein mistakenly thinks one can use this notion to *define* knowledge.
27. For an approach along these lines, see Craig (1990).
28. I wish to thank Chris Daly, as well as audiences at the University of Arizona, University of Rochester, and UC San Diego, for many helpful comments on the manuscript.

References

Bealer, George (1992). 'The Incoherence of Empiricism', *Proceedings of the Aristotelian Society* 66, supplement: 99–138.

Boghossian, Paul (1996). 'Analyticity Reconsidered', *Noûs* 30: 360–391.

BonJour, Laurence (1998). *In Defense of Pure Reason*. Cambridge: Cambridge University Press.

Craig, Edward (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon.

Frege, Gottlob (1980). *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*, second edition, tr. J.L. Austin. Evanston, Ill: Northwestern University Press.

Gardenfors, Peter (2000). *Conceptual Spaces: The Geometry of Thought*. Cambridge, Mass.: MIT Press.

Gettier, Edmund (1963). 'Is Justified True Belief Knowledge?' *Analysis* 23: 121–123.
Goddard, Cliff (2001). 'Lexico-Semantic Universals: A Critical Overview', *Linguistic Typology* 5: 1–65.
Goodman, Nelson (1955). *Fact, Fiction, and Forecast*. Cambridge, Mass.: Harvard University Press.
Huemer, Michael (2005). *Ethical Intuitionism*. New York: Palgrave Macmillan.
Hume, David [1758] (1975). *Enquiry Concerning Human Understanding*. Reprinted in *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, edited by P. H. Nidditch. Oxford: Clarendon.
Jackson, Frank (1994). 'Armchair Metaphysics', in Michaelis Michael and John O'Leary Hawthorne (eds), *Philosophy in Mind: The Place of Philosophy in the Study of Mind*. Dordrecht: Kluwer, pp. 23–42.
Kant, Immanuel (1965). *Critique of Pure Reason*, tr. Norman Kemp Smith. New York: St. Martin's Press.
Klein, Peter (1971). 'A Proposed Definition of Propositional Knowledge', *Journal of Philosophy* 68: 471–482.
Laurence, Stephen and Eric Margolis (2003). 'Concepts and Conceptual Analysis', *Philosophy and Phenomenological Research* 67: 253–282.
Locke, John [1690] (1975). *Essay Concerning Human Understanding*, edited by P. H. Nidditch. Oxford: Clarendon.
Oddie, Graham (2005). *Value, Reality, and Desire*. Oxford: Clarendon.
Price, H. H. (1969). *Thinking and Experience*, second edition. London: Hutchinson & Co.
Quine, Willard van Orman (1951). 'Two Dogmas of Empiricism', *Philosophical Review* 60: 20–43.
Rey, Georges (2004). 'The Rashness of Traditional Rationalism and Empiricism', *Canadian Journal of Philosophy* 34 (Supplement): 227–258.
Rodriguez-Pereyra, Gonzalo (2002). *Resemblance Nominalism: A Solution to the Problem of Universals*. Oxford: Clarendon.
Shope, Robert K. (1983). *The Analysis of Knowledge: A Decade of Research*. Princeton, N.J.: Princeton University Press.
Tkalčič, Marko and Jurij F. Tasič (2003). 'Color Spaces – Perceptual, Historical and Applicational Background', *EUROCON 2003: Computer as a Tool. The IEEE Region 8*, vol. 1, pp. 304–308, <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1248032&isnumber=27949>, accessed July 25, 2012.
Weinberg, Jonathan M., Shaun Nichols, and Stephen Stich (2001). 'Normativity and Epistemic Intuitions', *Philosophical Topics* 29: 429–460.
Williamson, Timothy (1995). 'Is Knowing a State of Mind?' *Mind* 104: 533–565.
Wilson, Mark (1982). 'Predicate Meets Property,' *Philosophical Review* 91: 549–589.

10.1057/9781137344557 - The Palgrave Handbook of Philosophical Methods, Edited by Christopher Daly