MOLYNEUX'S QUESTION AND COGNITIVE IMPENETRABILITY

BY JOHN CAMPBELL

1. Common Sensibles

Our understanding of the concepts we use in ordinary thought and talk often depends on our experience of the world. For instance, our understanding of words for the colours of objects seems to depend on our having had experience of the colours. It is a datum that someone who has never had colour experience will not know what the various colours are. Similarly, the ways in which we ordinarily understand words for the shapes objects commonly have, such as 'sphere' or 'cube', seem to depend on our having experienced shapes.

A common picture of the structure of our cognitive lives makes a distinction between the level of conceptual thought – the use of concepts in ordinary thought and talk – and the level of information-processing content described by scientists studying the brain systems involved in perception. Much discussion of cognitive architecture has to do with the relations between these two levels of content, conceptual vs. information-processing, and the internal structure of information-processing content in particular.

In this paper I want to propose that a full account of the structure of our cognitive lives will have to find a way of characterising our conscious experience of the world. For we cannot give an analysis of conceptual thought that misses out its relations to

experience, and we cannot give a characterisation of experience without considering its relations to information-processing in the brain.

A simple illustration of the need for a characterisation of the phenomenal level is provided, I will argue, by Molyneux's Question. This question concerns a man born blind, who has learnt by the use of his touch to distinguish spheres from cubes. We suppose him now to gain the use of his sight. And suppose a sphere and a cube to be placed where he can see but not touch them. Will he be able to tell which is the sphere, and which the cube? (For a survey of the history of this problem, see Morgan 1977.)

We can put the problem raised here as follows. Do we have just one repertoire of shape concepts, concepts such as 'sphere' or 'cube', and related geometric notions, that we apply indifferently on the basis of sight and touch? Or are there different concepts of shape, special to the various senses? In addressing this issue it is a natural thought that the shape of an object has significance for how it will behave in its interactions with other objects and with you. And the causal significance of possession of a particular shape will be the same, whether the shape is perceived through vision or touch. So perhaps our understanding of the causal significance of shape is the same, whether in vision or touch. Will that be enough to establish the sameness of the shape concepts we use on the basis of vision and touch?

To pursue this line of thought, we have to say something to explain in what sense the ordinary subject knows the causal significance of shape properties. On the face of it, you might argue that grasp of shape properties involves grasp of a theory articulating their significance. And this theory, you might say, penetrates both touch and vision. So you would argue that our ordinary shape concepts penetrate our perceptual systems, so

that once the newly sighted subject is enjoying ordinary vision, the ordinary shape concepts that he uses on the basis of touch will also figure in the content of his vision. In that case there would be no problem about him identifying the sphere and the cube, once the shape concepts had penetrated his visual system.

Alternatively, you might argue that visual shape perception is not cognitively penetrable by ordinary concepts, and that the right way to address the issue is to look at the relation between the conceptual level and the level of perceptual information-processing. In that case, you will take the question to be whether we have the same kind of information-processing content relating to shape in touch as in vision. And pursuing this line, you might say that grasp of the causal significance of shapes consists in our ability to use shape information in the control of action on objects, so that reaching and grasping movements, for example, are executed effectively.

In this paper I want to argue that setting up the issue in this way misses an important aspect of our understanding of shape concepts, namely their relation to our conscious experience of shapes. The key issues we have to address are whether conscious attention to the shapes of objects is amodal, or rather modality-specific; and we have to ask whether the ability to manipulate shapes to bring about changes in the behaviour of objects should be thought of as amodal, or is it rather modality-specific. Before sketching this way of formulating the issue, though, I will look further at the first two options: the idea that our concepts of common sensibles, such as shapes, penetrate vision and touch, and the idea that the key issue is whether information-processing content relating to shape is the same in vision as in touch. I begin with the idea that shape concepts penetrate vision and touch.

2. Common Sensibles as Depending on Commonsense Theory

We can see the shapes of things, whether they are round or oval and so on. We can feel the shapes of things – you can tell by touch whether something is round or oval. Why does it seem so evident that the shape properties we can see are the same properties as the shape properties we can touch? Why does it seem so obvious that seen roundness is the same property as felt roundness? Why, that is, is roundness a common sensible?

The first line of thought I want to consider runs as follows. The shape of an object has many implications for how it will behave in its interactions with other objects. There is much regularity governing the behaviour of something which is round, for instance. If it is also made of a rigid material, it will roll. If it is rigid, it will not stack together with other rigid round things without leaving gaps. And so on. The ordinary subject who can see the shapes of things knows that there are these regularities governing the behaviour of things with particular shapes; this is a taken-for-granted part of everyday thinking, used in the explanations we give and the predictions we make.

The regularities governing the behaviour of something seen to be round also govern the behaviour of something felt to be round. In our ordinary thinking, we take it for granted that the same regularities govern things seen to be round as govern the behaviour of things felt to be round. And, you might argue, that sameness of the regularities known by the subject to apply to seen and felt roundness is what constitutes

the sameness of the concept of roundness, whether it is used on the basis of vision or on the basis of touch.

You could put the idea by saying that we ordinarily grasp a naïve physics of our surroundings; we are able to explain and predict what is going on around us. We say, 'the peg could not get through the hole because it was square', 'the wheel could not take the weight of the cycle any more because it had buckled', 'the chairs could not be stacked together because they were all of different shapes', and so on. It is not that we explicitly articulate the generalisations we use here. Rather, it is just that we systematically do give such explanations and we are systematically able to make predictions about how the shapes of things will affect their behaviours. The knowledge of regularities here is 'implicit', not in the sense that it issues only in non-verbal behaviour, but rather in that it informs the explicit explanations and predictions we make, without itself being explicitly articulated.

In physics there are concepts such as 'neutrino' which you cannot understand unless you know something of the theory which uses them. You might suggest that there are such concepts in our naïve physics too, and that shape concepts are of this kind. To understand the concept, you must have this implicit grasp of the naïve theory which uses it. To understand a shape concept you must be able to use it in giving such explanations and making such predictions.

Evans presented a version of this idea in 'Things Without the Mind'. He said that to grasp the concepts of properties such as shape, 'one must master a set of interconnected principles which make up an elementary theory – of primitive mechanics – into which these properties fit, and which alone gives them sense.' (Evans 1985a, p.

269). The shape concepts are used in this theory in exactly the same way, whether they are being used on the basis of vision or touch.

Evans' idea here seems to have been that it is sameness of these 'interconnected principles' governing visual shape and tactual shape that constitutes the manifest sameness of the shape properties seen and touched. For as we shall see in a moment, he argues that without such a background in naïve theory, our concepts would be modality-specific. This is a matter of ordinary belief and reasoning penetrating vision and touch: it is a matter of how a commonsense grasp of the mechanics of our surroundings makes it possible for us to recognise the sameness of seen and felt shapes. If we did not grasp this mechanics, if we had only specifically visual or tactual processing of shape, so far as possible, in the different sensory modalities, there could be no recognition of the sameness of the properties experienced in different modalities. Such concepts of shape as we had could only be modality-specific.

On this analysis, it is because our thinking about colour properties, for example, does not embed them in any such theory, that colour concepts have to be specific to vision. So long as there is a commonsense theory embedding a concept, that theory can be held constant as the concept is used to classify sensory information in various different sensory modalities. But without the theory, says Evans, so long as we have only the perceptual processing, 'no single *sensory* property can be defined in relation to different senses.' (Evans 1985a, p. 270). This is why colours can only be seen, not recognised by means of any other sensory modality.

Having stated this view of the relation between commonsense theorising and shape perception, there is an alternative which it is natural to oppose to it (indeed, in his

later essay, 'Molyneux's Question', Evans presented a form of this alternative view, which seems flatly opposed to his earlier account). On this rival view, sameness or difference of shape across the sensory modalities may indeed be apparent to the subject; it may be apparent to the subject that it is the very same properties that are being seen and touched. But on this rival view, what is responsible for that is not sameness of the theory used by the subject. What is responsible is, rather, a more primitive similarity in the perceptual information provided by the brain mechanisms used in vision and touch. A sameness across the information-processing contents is what makes evident to the subject the sameness of the seen and felt properties. So on this kind of view, the existence of common sensibles, such as shape properties, does not depend on any cognitive penetration of vision and touch by commonsense theorising.

3. Cognitive Penetrability

I have set out the account on which the similarity between shape perception in touch and in vision is explained by the fact that touch and vision are both using shape concepts which have their meaning in virtue of their role in a single naïve theory which perceivers implicitly know. The problem for this view is that it threatens to make shape perception responsive to our common-sense reasoning. Of course, some aspects of perception are responsive to our common-sense reasoning. Suppose, for instance, that you see someone walking up your driveway who looks threatening. When you realise that this is the postman, the visual impression of threat may vanish entirely. The problem is that shape

perception does not seem to be of this sort. Illusions of shape, for instance, persist even though you know them to be illusions.

In Pylyshyn's terminology (Pylyshyn 1984, 1999), shape perception seems to be part of early vision, and early vision is not cognitively penetrable. Let me set out the issue here more fully, in Pylyshyn's terms. He gives a functional characterisation of 'early vision', saying that it 'involves the computation of most specifically visual properties, including 3D shape descriptions' (Pylyshyn 1999, p. 343). Early vision is governed by its own distinctive principles of visual computation. Pylyshyn points to two ways in which these principles governing early vision differ from the principles of rationality operative in ordinary thinking. First, the principles governing visual processing are defined over specifically visual primitives rather than physical properties of the seen objects. Thus there are labels for luminance, perceived relative size, and so on. The system has constraints on how these labels are to be applied over the entire scene. And the principles governing the labelling of a particular scene will be enforced by the visual system, even if the result is the illusion of an impossible scene (for example as in Escher prints). Secondly, the principles used by the visual system are not like principles of rationality in that they do not seem to find the simplest or likeliest interpretation of the scene, even given the information available to vision. Pylyshyn gives the example of 'amodal completion' – cases in which partially occluded figures are seen as whole figures partially hidden from view, rather than as fragmented figures which are entirely in view. (This is 'completion' because the visual system attempts to recover what the whole of the hidden figure is. It is 'amodal' in that the completion does not take the form of providing a visual impression of the hidden parts of the figure; the occluded

components remain occluded.) There are many cases in which the visual system will construct complex and unlikely occluded figures, following its own principles, even though the visual evidence allows of simpler and more likely interpretations. The point about the principles vision uses here is that they are 'natural constraints' which will typically work effectively to find the shapes of seen objects, given the kind of physical environment we occupy.

Early vision, then, recovers the 3D shapes of objects, using its own principles of computation. These principles are not the same as the general-purpose principles of rationality used in ordinary thinking. And Pylyshyn's thesis is that the computations performed in vision are 'cognitively impenetrable', in the sense that they are not rationally affected by the beliefs, desires or reasoning of the subject. Of course, the objectives of the subject may affect the direction of attention in early vision. But Pylyshyn's thesis allows for that. The point is that visual processing is not responsive to the beliefs and objectives of the perceiver in the same way as is any particular belief the subject has; it cannot be viewed as a matter of finding what is likely to be the case given the evidence available to the subject.

Here I do not want to attempt an evaluation of Pylyshyhn's thesis, but only to ask whether it is consistent with the view that the similarity between shape perception in touch and in vision is explained by the fact that touch and vision are both using shape concepts which have their meaning in virtue of their role in a single naïve theory which perceivers implicitly know.

A proponent of the view that the representation of shape in vision is conceptual might argue as follows. You might acknowledge that the visual representation of shape

is ‘belief–independent’ in the sense that visual illusions of shape may persist even though the perceiver knows them to be illusions. You might acknowledge that visual content is belief-independent and governed by its own specifically visual rules of processing. Nonetheless, you might maintain that the content of vision involves shape concepts which have their contents in virtue of their role in a common-sense theory. The picture is that the concept has its meaning in virtue of its role in a common-sense theory, but that nonetheless it is as it were hijacked by vision and used within vision subject to alien rules. Similarly, the concept proper to common-sense theory may be hijacked by touch and used there subject to principles specific to touch. The deliverances of vision and touch, using the concept in this way, will be belief-independent. But whether the subject accepts the testimony of the senses will depend on reasoning using the common-sense theory, in the light of all the available evidence.

The trouble with this picture is that it is not obvious that it makes sense to suppose that we have the very same shape concepts being used in vision as in ordinary reasoning, but that in vision those very concepts are governed by quite different rules to those of our common-sense theory. The position we are considering begins with the idea that the shape concepts have their meanings only in virtue of their roles in a common-sense theory. How then can those concepts retain those meanings when they are extracted from the theory and used in accordance with specifically visual or tactual principles that are quite different to those of common-sense theoretical reasoning?

You might draw an analogy between vision and the testimony of another person. When someone else tells you the shape of an object, the content of the testimony is certainly conceptual. And you can understand what the other person is telling you even if

you do not believe it; so testimony is in that sense 'belief-independent'. But this model of testimony is quite misleading here. It is true that testimony is belief-independent, rather as perceptual content is belief-independent. But the person who is giving you the information is using the shape concept in accordance with the common-sense theory, even if in their speech they mislead you, so the case is not analogous to that of the visual system. Of course, you might say that the individuation of shape concepts has nothing to do with their roles in a theory, so there is no problem in using one and the same concept now in the context of a theory, now in the quite alien context of visual processing. If you take that view, though, you lose the right to say that what unites shape concepts across the sensory modalities is that they have their meanings in virtue of their embedment within a single theory.

4. Sameness of Information-Processing Content as Grounding Common Sensibles

You might draw the moral from the above discussion that the content of perception is not itself conceptual, but is prior to conceptual content. You might then hold whether shape concepts can be applied indifferently on the basis of vision and touch depends on a prior issue: whether, at the information-processing level, vision and touch represent shapes in just the same ways. If vision and touch represent shapes in different ways, then the newly sighted subject will not be able to tell, by vision alone, which is the sphere and which the cube; the shape concepts he already has will be specifically tactual. If, on the other hand, vision and touch do, at the information-processing level, represent shapes in

the very same ways, the newly sighted man will be able to use vision to apply the concepts he has already formed on the basis of touch. For, you might argue, what he has learned, through his training with touch, is how to apply shape concepts on the basis of information-processing content. So if information-processing content of the very same sort that is available in touch becomes available in vision, then the newly sighted man will simply be going on as before when he applies the shape concepts on the basis of vision.

On this approach, the question is whether the neural systems involved in visual and tactual information-processing are representing shapes in the same ways. The idea here is that without appealing to facts about concepts, we can establish that the visual and tactual information-processing systems represent shapes in the same ways. Then we can conclude from this that the shape concepts we have are applied indifferently on the basis of vision and touch.

You might argue as follows: visual and tactual shape perception play the very same roles in controlling our actions; since they play the same roles in controlling our actions, they must be representing the shapes in the same ways. I will call this, the Action Argument:

> The Action Argument: visual and tactual shape perception play the very same roles in controlling our actions; since they play the same roles in controlling our actions, they must be representing the shapes in the same ways.

This Action Argument is a cleaned-up version of the main argument that Evans gives in his classic discussion, 'Molyneux's Question'. Evans' argument there is quite elaborate. Strictly speaking, he does not appeal to the idea that there is any processing of shape information in vision or touch. His argument appeals to contents relating to shape only at the conceptualisation stage, when the subject applies shape concepts on the basis of visual or tactual perception. According to Evans, shape concepts are applied on the basis of information about the egocentric locations of the things around the subject. And his argument is that we are dealing with the same egocentric space in every sensory modality, because all egocentric information affects the very same repertoire of behaviours: 'there is only one egocentric space, because there is only one behavioural space' (Evans 1985b, 389-390). Now the issues about egocentric space are of interest in their own right. But I set them aside here, because there seems to be little reason to accept Evans' idea that information about egocentric location is the basis on which we apply shape concepts. There is ample evidence for the existence of form processing in early vision – indeed, as we saw, Pylyshyn takes shape processing to be the principal task of early vision. The shape of the object is established by the visual system, prior to conceptualisation. It is not credible that the subject making a conceptual judgement as to the shape of the object does not use this visual information about its shape, but prefers instead to rely on an inference from the egocentric locations of the parts of the object, as Evans supposes (Evans 1985b, p. 389). The Action Argument recasts Evans' reasoning to accommodate this point. So long as shape perception affects behaviour in the very same way whether in vision or in touch, Evans can recast his argument so that it dispenses with the shuffle through egocentric space.

The fundamental problem for the Action Argument is its uncritical use of the notion of ‘sameness of representation’ or ‘sameness of representational content’. We are considering vision and touch as distinct information-processing systems, each with its own domain of input data and its own range of computations to perform on that data. What sense does it make to talk of ‘sameness of representation’ or ‘sameness of informational content’ across input systems? Of course, we as theorists can, if we like, introduce ways of comparing representations in different input systems. For instance, you might say that two neural patterns, in two different input systems, carry the same information just if the occurrence of those two neural patterns is, in ordinary cases, produced by just the same external stimuli. With some work, you might construct a reasonably well-defined notion of ‘sameness of informational content’ along those lines, which would allow you to compare representations in different input systems for sameness or difference of content. But in any particular case, there would be no guarantee that the sameness or difference of informational content, so defined would be recognised by the informational system at any point. Suppose that the two input systems are both input to some third system. Whether that third system treats in the same way two representations that by your definition ‘have the same content’ is a further, empirical question. The third system might treat in the same way two representations that by your definition ‘have different contents’. Or it might process in quite different ways two representations, from the two input systems, that by your definition, ‘have the same contents’.

Indeed, at the information-processing level, there is no point in asking about sameness or difference in content, in the abstract, for contents in different input systems.

The only question it makes sense to ask is whether the contents are treated in the same ways by whatever system or systems they are output to. The mistake is to suppose that there is a prior notion of sameness of content to which we can appeal, so that we can conclude that sameness of content in this prior sense means that any system to which the two contents are output must, or at any rate ought to, treat the two contents in the same way.

The Action Argument claims that visual and tactual shape perceptions have just the same impacts on the control of action, and concludes from this that visual and tactual shape perceptions have just the same impacts on the application of concepts to objects. But the conclusion does not follow from the claim. It is entirely possible that two representations might be treated in the same way for the purposes of control of action, but treated quite differently when it comes to the application of shape concepts to physical objects.

In fact, the situation for the Action Argument may be even worse than I have suggested. Goodale and Milner 1995, and Jeannerod 1997, among many others, have argued that we must distinguish between two visual pathways, one used for the control of action and the other used for identification and recognition. If that is correct, then it may be that the shape representations used for the visual control of action are not the same as the shape representations used for the explicit classification of shapes by the subject. So a similarity between the shape representations used for the visual control of action and the shape representations used for the tactual control of action would not imply a similarity between the shape representations grounding the application of shape concepts

on the basis of vision, and the shape representations grounding the application of shape concepts on the basis of touch.

The problem that is raised by our discussion to this point is how we are even to make sense of the idea that it might be the same shape concepts that are being used in vision as in touch, Saying that it is the same shape concept if it is embedded in the same theory seemed inconsistent with the cognitive impenetrability of early vision. But saying that it is the same shape concept if it is responsive to the same perceptual information as input appeals to a notion of sameness of perceptual information across different input systems that does not seem to be well-defined for our purposes; we do not know whether we have the same perceptual information output from vision and touch to the application of concepts until we know whether we have the same concepts being used in sight as in touch.

5. Conscious Attention to Shape for Manipulation

On a classical approach to semantic theory, grasp of the concept expressed by a predicate, such as a shape predicate, is a matter of knowing which property the concept stands for. I have in effect been taking it throughout this essay that knowing which property a shape concept stands for is a matter of knowing the causal significance of the shape property; that properties are individuated by their causal roles. The question then is whether knowledge of the causal significance of a shape property should be thought of as amodal. The fundamental problem here is how to characterise the format in which the ordinary

subject has knowledge of the causal significance of the shape property. The most obvious proposal is that the subject has explicit theoretical knowledge of the causal role of the property; but we have set aside that idea, as implying that there can be perception of shape only insofar as there is cognitive penetration of vision.

The idea of appealing to a prior notion of sameness of perceptual content in early vision and in touch, that we have just been considering, attempts to finesse the question of the format in which the subject knows the causal significance of shape properties. The implicit suggestion is that we have a grasp of the notion of sameness of information that is prior to any issue about the format in which specifically causal information is held, and that this prior notion of sameness of information can be used to explain how the subject can recognise sameness of shape across modalities. As we have seen that approach does not seem to help, as we have no such prior notion of sameness of perceptual information. Sameness of implication for action is one thing, and recognition of sameness by the subject is another.

I think, though, that there is another way in which we can frame the problem raised by Molyneux's Question, and I want to end by sketching this different perspective. I want to propose that we can think of knowledge of which property a shape concept stands for as being provided by the capacity for a particular kind of conscious attention to the shapes of things. And the key idea underlying a positive answer to Molyneux's Question is this: if there is at work a single capacity for conscious attention to a particular aspect of the world, then it must be apparent to the subject that it is a single aspect of the world that is in question.

Let me give an analogy. Suppose that a subject has a particular way of classifying the things he sees, so that it usually makes sense for him to ask how many things of that sort he is currently seeing. And suppose that this subject also has a way of classifying the things he hears, so that it usually makes sense for him to ask how many things of that sort he is currently hearing. Suppose now that this subject is able to attend consciously to the number of things (of the given sort) that he is seeing; he can exercise this capacity for conscious attention by counting. Suppose also that our subject is able to attend consciously to the number of the things (of the relevant sort) that he is hearing. Of course, the classificatory abilities that the subject has may each be modality-specific. But further than that, we can ask whether the ability to count, as such, is modality-specific. That is, we can ask whether the ability to attend consciously to the number of things presented is an amodal skill, applied indifferent in each of the two sensory modalities. Or is this ability to count specific to each modality, so that the ability is learned and exercised independently in sight and in hearing? Now if counting is a genuinely amodal skill, as we would naturally suppose it to be, we would also expect it to be apparent to the subject that the numbers he assigns to the things he can see are just the same as the numbers he assigns to the things he can hear. He should take it for granted that it makes sense, for example, to add together the number assigned to the things he sees and the number assigned to the things he hears. So we take it for granted that sameness of the phenomena to which he is consciously attending should be apparent to the subject. I think that this is an aspect specifically of our conscious life; it is an aspect of our cognitive structure that relates specifically to the phenomena to which we consciously attend. In effect, I have in the last section been arguing that there is no such architectural

constraint on the ascription of content to information-processing in the visual pathways of the brain.

Let me set out how this approach bears on Molyneux's Question. In the early phase, the Molyneux subject has learned the capacity for conscious attention to the shapes of the objects he touches. There is more to this than merely being able to identify the shapes of objects as instrumental to, for instance, semantic classification of them. Classifying the thing before you as, say, a particular type of animal, may well involve, as instrumental to being able to effect the classification, that the visual system finds the shape of the object. Similarly, if you are to read a sentence, your visual system must have found the shapes involved; but that is not yet for you to have consciously attended to the shape of the sentence.

One task for which conscious attention to the shape of an object is needed is, evidently, application of the right shape concept to the object. But I want to focus on another use of conscious attention to shape. This is the ability to manipulate the shape of an object to produce a desired effect. For instance, you might squash a package to make it go through a letterbox, adjust the shape of a cushion to make a chair comfortable, or put together building blocks in a new configuration so that you have a stable structure. Possession of this second ability seems more informative about your possession of a shape concept than does the mere ability to produce a label for a perceived shape on demand. Grasp of concepts of shape is not merely a matter of being able to tell when the concepts apply to one object or another. You have to know something of the significance of the object being a particular shape. And that grasp of significance is constituted by your practical knowledge of how to manipulate the further characteristics of the object –

whether it will go through the letterbox, whether it will be comfortable to sit on, whether it will be a stable structure, but manipulating its shape. And that capacity to manipulate the shape of the object demands conscious attention to the shape of the object. It goes far beyond the mere use of shape information to control the prehension of the hand, for instance. It requires that the subject have the idea of the shape of the object as a variable whose manipulation will affect the further characteristics of the object.

There is a contrast here between an understanding of shape concepts and an understanding of colour concepts. The colours of objects are often symptomatic of their further characteristics. For example, the colour of a food is often a good guide to whether it is ready to eat, fresh, and so on. But in general, you cannot manipulate those further characteristics of the object by manipulating its colour. You cannot make the food fresh, for instance, by lightening its colour.

I think that this point gets at what is right in Evans' idea that shape concepts have their meaning in virtue of their roles in an implicitly grasped mechanical theory. Evans' characterisation did not manage to articulate the sense in which this theory is a causal theory. Demanding that the subject should have the ability to use shape as a variable whose manipulation allows the manipulation of further characteristics of the object does do something to explain the sense in which the subject can be said to grasp the causal significance of shape properties.

The problem that was raised by the discussion above, of the idea of shape concepts as embedded in a primitive theory, and the idea of shape concepts as responsive to a more primitive level of perceptual information, was how we are even to make sense of the possibility that it is the same shape concepts that are being used in relation to the

different senses. But now I want to propose that we can think of Molyneux' Question as raising two problems which seem eminently amenable to empirical study.

First, there is the question whether the capacity for conscious attention to the shapes of objects is amodal. Is there a single capacity, the ability to attend to the shape of a perceived object, which is being exercised indifferently in sight and touch? Or is it that attention to shape in vision and attention to shape in touch are separate abilities, which may be acquired and exercised independently of one another? (Of course, these two extremes do not exhaust the possibilities. There are parallels between this problem and the question whether the capacity to attend to locations is amodal; for a review of some of the issues raised by that question see Spence et. al. 2000.) Secondly, there is the question whether the ability to manipulate the shapes of objects to affect their further characteristics is amodal, or whether again it is acquired and exercised independently in vision and in touch. (Again, these are two extreme positions and there is a family of further distinctions among possible positions to draw here.)

These two questions are related. When you attend specifically to one aspect or another of the scene before you, it always makes sense to ask: what is the task in whose service you are attending to that aspect of the scene? What I am asking could be put as a single complex question: is the ability to attend consciously to the shapes of objects, for the purpose of further manipulation of their characteristics, a single attentional skill, no matter the modality in which it is exercised? And I am saying that if it is a single capacity for conscious attention that we have here, then we would expect it to be apparent to the subject that it is the very same shape properties that he attending to on the basis of vision as he attends to on the basis of touch. This is a point about the architecture of

specifically conscious attention; we simply miss this point if we look only at the structure of information-processing in the brain. And this point – an aspect of the unity of consciousness – underlies the architecture of our conceptual thought.

REFERENCES

Evans, Gareth. 1985a. 'Things Without the Mind'. In his *Collected Papers*. Oxford: Oxford University Press.

Evans 1985b, 'Molyneux's Question' In his *Collected Papers*. Oxford: Oxford University Press.

Goodale, Melvyn A. and Milner, A.D. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.

Jeannerod, Marc. 1997. *The Cognitive Neuroscience of Action*. Oxford: Blackwell.

Morgan, Michael J. 1977. *Molyneux's Question: Vision, Touch and the Philosophy of Perception*. Cambridge: Cambridge University Press.

Pylyshyn, Zenon. 1984. *Computation and Cognition*, Cambridge, Mass.: MIT Press.

Pylyshyn, Zenon. 1999. 'Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception' *Behavioral and Brain Sciences* 22, 341-423.

Spence, C,., Pavani F and Driver J. 2000. ‘Crossmodal Links between Vision and Touch in Covert Endogenous Spatial Attention’. *Journal of Experimental Psychology: Human Perception and Performance*, 26, 1298-319