Objects, affordances …action!

At the 2000 London Conference **GLYN HUMPHREYS** *gave his Presidents' Award Lecture on the cognitive neuroscience of action selection.*

IN everyday life we carry out many hundreds of visually guided actions on the objects that surround us. We may reach and grasp a kettle and pour boiling water from it into a teapot; we may pour from a jug of milk and from a teapot into a cup; and we may raise the cup to our lips to drink. Although each of these actions seems simple enough, the processes involved are complex. We need to use visual information to guide the reach-and-grasp actions. Having grasped an object, we must then effect the appropriate category of action – we need to pour from the jug but drink from the cup.

Over the past 20 years a considerable amount of research has been carried out into understanding the first part of the process, how visual information is used in reach-and-grasp actions (see Milner & Goodale, 1995, for one review). However, much less work has been conducted into the factors that determine the categories of action we perform with objects, once the objects have been grasped.

In this article I will discuss research within my laboratory that has focused on this last question – how categories of action are selected from visually presented objects. The research uses converging evidence from experimental psychology, cognitive neuropsychology and functional brain imaging to reveal both the underlying cognitive architecture and the neural substrate of action selection. It then uses computational modelling to provide an explicit account of both normal performance and its breakdown following brain damage. The work suggests that, in addition to being based on contextual and associative knowledge about objects, action selection is influenced by 'affordances' derived from the visual properties of objects.

For example you might use the object shown in Figure 1 appropriately because you remember encountering it in a kitchen, that it was used with an olive, and so on (i.e. using associative and contextual knowledge). Alternatively, you might use it correctly, even without having encountered such an object before, because its structure indicates directly the possible actions that may be performed (i.e. you are influenced by the object's 'affordance' for action).

There are important implications of this idea for understanding how behaviour must be controlled when performing everyday actions. I will discuss these implications in relation to work in both cognitive neuropsychology and developmental psychology.

Neuropsychological evidence for affordances

The term *affordance* was introduced into the literature by the ecological psychologist J.J. Gibson (1979) to describe the possibilities for action offered by objects. Gibson pointed out that, depending on our current behavioural goal, the same object could 'afford' different actions – a cup may be used for drinking, but it could also be used to catch a spider if that was what was needed. Gibson claimed that objects were perceived in terms of these affordances for action. Although the idea has had its advocates (e.g. Bruce *et al.*, 1996), there has been remarkably little supportive evidence that affordances mediate action independently of semantic (contextual and associative) influences.

Part of the problem in discovering whether affordances mediate action has been in defining what is meant by the term *affordance*. What I mean by the term is some direct link between the perceived visual properties of an object and an action that may be performed with it. For example, the joint holding the top arm of the object in Figure 1 constrains the arm so that it may only be pushed down, and the hole in the lower part of the object can be filled by the vertical section of the top arm, and so forth; these visual properties may directly signal the action of pushing the top arm down. This link between the perceived properties of objects and actions may be, but does not need to, be based on a stored representation of the particular object.

Crucial to the idea of affordances is that, as mentioned above, particular categories of action may be cued directly by the visual properties of an object even if the object has never been encountered before. Although many people will never have seen an object such as that illustrated in the Figure 1, there is nevertheless high consistency in judgements of how it might be used ('push the top left part down!').

FIGURE 1 An instrument to remove a stone from an olive. Although this object is very unfamiliar to many observers, there is nevertheless high agreement about how it might be used

I first entertained the idea that affordances can play a role in action selection when conducting neuropsychological research. Jane Riddoch and I (Riddoch & Humphreys, 1987) worked with a patient JB, who was often unable to name objects correctly but who could nevertheless gesture how the objects should be used. The gestures could be highly specific, such as making a jabbing gesture to a fork with his left hand but a cutting gesture with his right hand to a knife. Yet, despite this, JB was impaired at judging which objects might be used together (e.g. matching together a knife and fork and rejecting a teaspoon), indicating a problem in accessing what might be termed 'semantic' memory (i.e. associative or contextual knowledge). To account for the results, we argued that JB could make gestures using direct, learned associations between objects and action. These gestures could be effected even when he could not gain access to detailed semantic knowledge. This suggests that there may be a direct 'route' to action not mediated by semantic knowledge: an 'affordance' route to action.

Jane Riddoch, Cathy Price and I (Riddoch *et al.*, 1989) also observed what in many respects appeared to be the opposite pattern of deficit, in the syndrome of 'visual apraxia' (see also DeRenzi *et al.*, 1982). The term apraxia is applied to patients who fail to use objects correctly even when they have no motor or recognition deficit. In visual apraxia, this problem is most pronounced when objects are presented visually, though the same patients may make appropriate gestures when given the object's name, and they may even be able to name the object in front of them! Here recognition is intact, but visually driven actions are impaired.

Such an impairment could occur if damage to the visual route to action 'blocks' action selection from other modalities. For instance, correct gesturing to an object name may involve retrieving semantic information about the object's function and using this information to select the appropriate action. In visual apraxia the damaged visual route seems to prevent the apparently intact semantic route from operating normally – perhaps because there is competition when multiple affordances are offered through the visual route. It is difficult to understand this pattern of behaviour unless a visual route could modulate semantic retrieval processes.

More recently I have provided other evidence for affordances directing action from the syndrome of unilateral neglect (Humphreys & Riddoch, 2001). This syndrome is characterised by patients failing to respond to stimuli presented on the side of space contralateral to their lesion. In one interesting case we found that a neglect patient could detect targets in multi-object displays when cued with an action ('find the object to drink from') but not when cued with the object's name ('find the cup'). This cueing-by-action was successful when the objects were oriented with their handles pointing towards the patient, but not otherwise. We suggested that a visual route to action was left relatively intact in this case and that the lesion disrupted the ability of the patient to use semantic cues to direct action. The 'action benefit' was found only when objects were placed in an appropriate orientation for a grasp response. This further suggests that the affordance of the object for action was important, rather than semantic cues (which would be the same irrespective of the orientation of the objects; indeed object orientation had no effect when the patient was cued by the target's name).

FIGURE 2 (a) Example of a non-object created by interchanging the locations of its parts, so that it can no longer be used in a plausible manner. (b) Example of a non-object created with parts in positions where they can easily be used

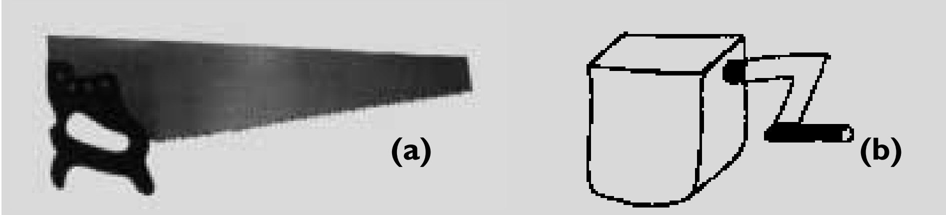

Evidence from experimental psychology

Neuropsychological cases can provide striking evidence for either the selective preservation or disruption of a visual route to action. Evidence from normal subjects, however, is less clear. Some anecdotal findings come from diary studies of errors made in everyday action. For example, Reason (1984) noted that participants sometimes reported making visual errors in which they mistakenly used one object as a visually similar neighbour (e.g. using hair spray as a deodorant). Such errors could arise because people make actions directly linked with the visual properties of objects, which overlap between the neighbour and the object for the intended action.

Raffaella Rumiati and I (Rumiati & Humphreys, 1998) provided evidence that similar errors could be elicited under laboratory conditions when 'normal' participants are forced to act to a fast response deadline. We had participants make gestures to pictures of objects, and asked them to try to complete the gestures before a tone was presented. The error rate for actions was increased under these conditions, with a relatively high proportion of the errors being visually related to the target (e.g. making a hammering gesture to the picture of a razor). Proportionately fewer errors were semantically related to targets (e.g. making a lathering gesture to the face when a razor was presented).

In contrast to the results with pictures, when the stimuli were presented as words only semantic errors arose. The semantic errors presumably reflect forced gestures to be made when semantic access to action was incomplete. Similarly, the visual errors may reflect activation when a visual route to action is not completed before the deadline. Since visual errors predominated with objects, it appears that there is early activation of actions based on partial visual cues, so that errors under deadline conditions reflect these actions rather than actions cued by semantic properties of the stimuli.

The nature of the direct route to action

Evidence from normal participants, then, converges with neuropsychological data in suggesting that categories of action can be cued directly by the visual properties of objects. Other neuropsychological research

Requests for reprints of this article should be addressed to:
Professor Glyn Humphreys, School of Psychology, University of Birmingham, Birmingham B15 2TT
Tel: 0121 414 4930
E-mail: g.w.humphreys@bham.ac.uk

FIGURE 3 The functional architecture of NAM (naming and action model), used to account for both normal and pathological patterns of action selection by Yoon *et al.* (in press)

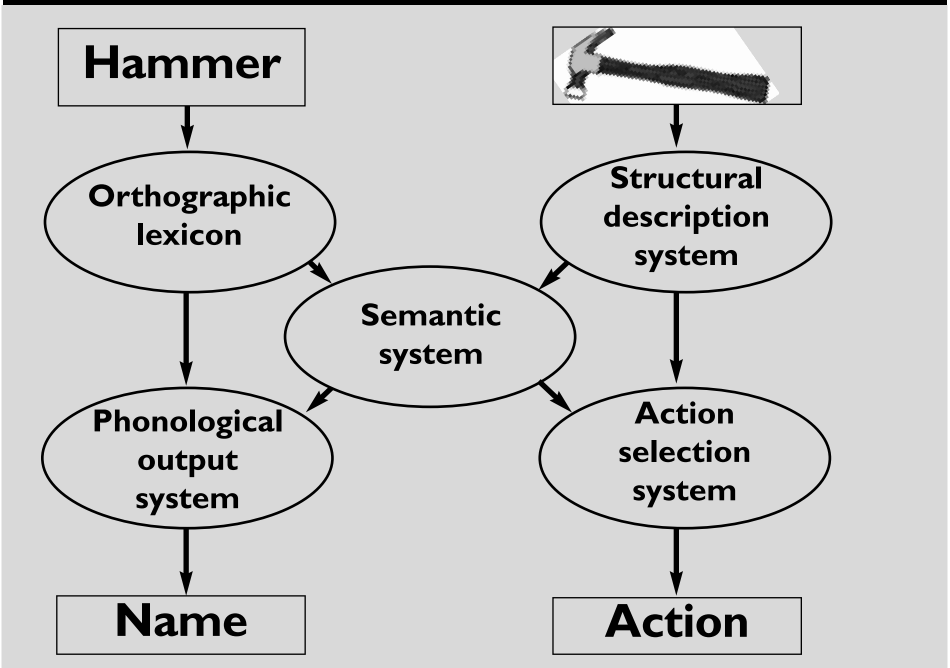

helps to reveal the nature of this direct visual route – and in particular, the nature of the visual information that supports responses through this route.

Marie Pierre Vernier, Jane Riddoch and I have worked with one patient, JP, who suffered a stroke affecting the left inferior frontal cortex. Following this, JP was selectively impaired at responding to tools and body parts. When given a name she was unable to point to a tool or body part shown amongst other items from the same category. She was also poor at matching these objects with associates (e.g. match a hammer with a nail, either presented as an object alongside the target object or as a name). This suggests that there was a problem in retrieving semantic information about these items.

We then presented JP with an 'object decision' task, in which she was presented with pictures and had to decide if they depicted real objects or nonsense objects ('non-objects'). She performed object decisions at a high level when the non-objects were created by altering the spatial relations between the parts of objects so that the objects could not easily be used (Figure 2a). However, she was quite impaired when the non-objects had plausible parts in the correct spatial relations (see Figure 2b).

The ability to carry out some types of object decision, in the face of very impaired semantic matching, suggests that JP could access some form of pre-semantic stored knowledge about objects. This process was sensitive to whether the parts of objects were positioned to enable the objects to be easily used, but it was not particularly sensitive to whether the stimulus as a whole was visually familiar. Hence JP was impaired at discriminating non-objects that were visually unfamiliar as a whole, but whose parts were in the correct position for object use. It seems that, in vision, objects such as tools are encoded on the basis of whether the parts are in the correct relative locations for action. A saw is encoded and recognised on the basis of whether it has a handle and a serrated edge for cutting, with these two parts in the appropriate spatial position for the object to be used in a sawing action (see Figure 2).

Interestingly, JP was sometimes able to gesture to objects, despite her semantic impairment. However, such gestures seemed very often to be plausible attempts at object use rather than indicating access to stored object–action associations. For instance, she used a razor to scrape along a table rather than to shave. This is just what we would expect if her actions were dependent on affordances, derived from the visual representations of objects, instead of being constrained by semantic knowledge.

Neural substrates of affordances?

A number of studies, using functional brain imaging, have demonstrated that, even in tasks such as object naming, there is activation of areas of the cortex associated with object use and movement (e.g., Chao *et al.*, 1999; Grabowski *et al.*, 1998; Grafton *et al.*, 1997; Martin *et al.*, 1995). These studies suggest that knowledge of object use may be contacted to facilitate the name retrieval process, for instance by enabling functionally different objects to be differentiated (see Humphreys *et al.*, in press). However, the studies are not informative about how this knowledge of object use is accessed – does access operate from semantic or visual representations?

To address this last question, Jackie Phillips, Cathy Price, Ute Noppeney and I conducted a study using PET methodology to image areas of cortex specifically activated in the retrieval of action information from objects relative to words (Phillips *et al.*, in press). Participants saw pictures of objects, non-objects or words and they had to decide whether a 'twist' or 'pour' action would normally be made to each target item. The brain activity involved in each of these tasks was then compared with that involved in a control task, where size judgements were made to the same stimuli. We found that the left occipito-temporal region was selectively activated by non-objects and objects in the action task compared with the control task, and the difference between the action and control tasks in this region was larger for non-objects and objects than for words.

Occipito-temporal regions of the brain have been associated with high-level processing of the structural properties of objects (e.g. Ungerleider & Haxby, 1994). The imaging data suggest that, when actions are made to objects, there is enhanced processing of the structural properties of the stimuli that might support these actions (with non-objects and objects, relative to words). Consequently increased activation is found in areas of cortex associated with processing of the structural properties of objects (in occipito-temporal cortex). This fits with our argument from patient JP, that structural processing may be sensitive to how parts are coded for action.

A computational account

The experimental and neuropsychological research I have outlined has highlighted that visual information can directly cue categories of action. This process can be preserved even when access to semantic knowledge breaks down (e.g. in the syndrome of optic aphasia). And, when visual links to action are damaged, retrieval of actions through a semantic route can be impaired (in visual apraxia). This last result indicates that the visual route to action can modulate semantic retrieval processes.

To provide a more detailed account of such interactions, it can be useful to develop explicit, computational models of performance, where interactions between processing routes have to be spelled out. One attempt to do this here was made by Eun Young Yoon, Dietmar Heinke and myself (Yoon *et al.*, in press). The model, known as NAM (naming and action model) uses the architecture shown in Figure 3. NAM adopts an 'energy minimisation' approach, in which responses are effected once activation within the network enters a stable state (a so-called 'basin of attraction').

In NAM the visual route provides a first 'push' within the network towards an appropriate basin of attraction for a given object, a process which is then supplemented by activation along the semantic pathway. However, since the visual route provides the first push, the network can continue to operate to generate action from an object even when the semantic route is damaged. This simulates the syndrome optic aphasia. The converse of this, however, is that damage to the visual route can push activation away from the appropriate basin of attraction, even though the semantic route is preserved. Then, performance can be worse when visual input is given than when there is no visual input and actions are derived solely from semantic/contextual knowledge. This simulates visual apraxia.

It may be noted that it is difficult to understand visual apraxia in more traditional cognitive models with independent 'dual-route' architectures (see e.g. Coltheart, 1985, for this approach to word recognition). This is because, in traditional dual-route architectures, an intact semantic route to action should operate in isolation from the visual route, and the semantic route appears intact in visual apraxia. In NAM, though, the visual route modulates semantic retrieval of action, even though the architecture of the routes enables each to be damaged selectively. It seems that interactions of the kind simulated by NAM are necessary to account for human neuropsychological data.

Implications of having a direct visual route to action

The approach I have outlined to the question of how categories of action are selected has emphasised the importance of drawing on converging evidence, in order to provide as full a picture as possible of the cognitive architecture underlying performance. Thus I have argued that experimental studies with normal participants, neuropsychological studies, studies using functional brain imaging and studies using explicit computational models can complement one another, demonstrating not only the existence of particular 'routes to action' but also something of their nature and how the routes interact. The functional imaging and neuropsychological research also speaks to issues of where these different routes may be localised within the brain.

If the argument for a direct, affordance-based route to action is correct, then there are important implications for understanding how actions are controlled – particularly when we enter environments containing multiple objects rather than the single objects used in many of the studies.

Take an everyday life task like making a cup of tea. Here we must use each individual object appropriately, and efficient task performance depends also on carrying the steps in the task in an optimal order, without repeating an action or omitting early actions in the sequence (e.g. forgetting to put water in the kettle). We must also ignore irrelevant objects that happen to be in the room. Now if some objects have stronger affordances than others, then actions may be strongly activated by those objects even when the objects should not be used at that particular time. Thus there will need to be some form of control process that itself modulates the effects of affordances on action selection.

A neuropsychological disorder, known as 'action disorganisation syndrome', has been identified in which patients fail to perform everyday life tasks in an organised manner. At least in part this appears to be because they are strongly cued by the environment to act on objects rather than acting within the constraints of the task (e.g. Humphreys & Forde, 1998; Schwartz, 1995). In some developmental pathologies, such as autism, there may also be strong propensities to act to 'pre-potent' environmental cues (Hughes *et al.*, 1994).

The idea of a direct, 'afforded' route to action provides one way to conceptualise aspects of these disorders. It is for future work to link models that specify how actions are activated from single objects with models that specify how we carry out tasks with multiple steps, where the appropriate objects have to be selected for action as well as the appropriate actions selected from objects (see Riddoch *et al.*, 2000a, 2000b). These latter models will have important implications for understanding many aspects of everyday behaviour and for understanding a variety of pathologies.

References

Bruce, V., Green, P. R., & Georgeson, M. A. (1996). *Visual perception: Physiology, psychology and ecology* (3rd ed.). London: Psychology Press.

Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience, 2*, 913–919.

Coltheart, M. (1985). The cognitive neuropsychology of reading. In M. I. Posner & O. S. M. Marin (Eds.), *Attention & performance XI*. Hillsdale, NJ: Lawrence Erlbaum.

DeRenzi, E., Faglioni, P., & Sorgato, P. (1982). Modality-specific and supramodal mechanisms of apraxia. *Brain, 105*, 301–312.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton-Mifflin.

Grabowski, T. J., Damasio, H., & Damasio, A. R. (1998). Premotor and prefrontal correlates of category-related lexical retrieval. *Neuroimage, 7*, 232–243.

Grafton, S. T., Fadiga, L., Arbib, M. A., & Rizzolatti, G. (1997). Premotor cortex activation during observation and naming of familiar tools. *Neuroimage, 6*, 231–236.

Hughes, C., Russell, J., & Robbins, T. (1994). Specific planning deficit in autism: Evidence of a central executive dysfunction. *Neuropsychologia, 32*, 477–492.

Humphreys, G. W., & Forde, E. M. E. (1998). Disorder action schema and action dysorganization syndrome. *Cognitive Neuropsychology, 15*, 771–811.

Humphreys, G. W., & Riddoch, M. J. (2001). Detection by action: Evidence for affordances in search in neglect. *Nature Neuroscience, 4*, 84–88.

Humphreys, G. W., Riddoch, M. J., & Forde, E. M. E. (in press). The principle of target-competitor differentiation in object recognition and naming (and its role in category effects in normality and pathology). In E. M. E. Forde & G. W. Humphreys (Eds.), *Category specificity in mind and brain*. London: Psychology Press.

Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science, 270*, 102–105.

Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. London: Academic Press.

Phillips, J., Humphreys, G. W., Noppeney, U., & Price, C. J. (in press). T*he neural substrates of action retrieval from objects and words: A PET analysis. Visual Cognition.*

Reason, J. T. (1984). Lapses of attention in everyday life. In W. Parasuraman & R. Davies (Eds.), *Varieties of attention*. Orlando, FL: Academic Press.

Riddoch, M. J., & Humphreys, G. W. (1987). Visual object processing in a case of optic aphasia: A case of semantic access agnosia. *Cognitive Neuropsychology, 4*, 131–185.

Riddoch, M. J., Humphreys, G. W., & Edwards, M. (2000a). Visual affordances and object selection. In S. Monsell & J. Driver (Eds.), *Attention & performance XVIII*. Cambridge, MA: MIT Press.

Riddoch, M. J., Humphreys, G. W., & Edwards, M. G. (2000b). Neuropsychological evidence distinguishing object selection from action (effector) selection. *Cognitive Neuropsychology, 17*, 547–562.

Riddoch, M. J., Humphreys, G. W., & Price, C.J. (1989). Routes to action: Evidence from apraxia. *Cognitive Neuropsychology, 6*, 437–454.

Rumiati, R. I., & Humphreys, G. W. (1998). Recognition by action: Dissociating visual and semantic routes to action in normal observers. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 631–647.

Schwartz, M. F. (1995). Re-examining the role of executive functions in routine action production. *Annals of the New York Academy of Sciences, 769*, 321–335.

Ungerleider, L. G., & Haxby, J. V. (1994). 'What' and 'where' in the human brain? *Current Opinion in Biology, 4*, 157–164.

Yoon, E. Y., Heinke, D., & Humphreys, G. W. (in press). Modelling direct perceptual constraints on action selection: The naming and action model (NAM). *Visual Cognition.*

SOME GOOD REASONS FOR WRITING FOR *THE PSYCHOLOGIST…*

Since our article appeared in The Psychologist *the media coverage our research group has received has rocketed. We have appeared on BBC TV and radio, independent TV and have already picked up two exciting new projects as a direct result of this exposure.*
Dr Jonathan Freeman, Goldsmiths College 'Immersive television', April 2001

Whereas writing papers with the RAE in mind is little fun and for a limited readership, writing for The Psychologist *is good practice for presenting to a much wider audience, who will actually read and maybe enjoy what you write.*
Professor Jim Horne, Loughborough University 'State of the art: Sleep', June 2001

I've never had such nice feedback on anything I've written before and it's still arriving.
Professor Jenny Firth-Cozens, University of Northumbria 'Why I study... stress in doctors', April 2001

My colleagues and I were delighted with our article. It created considerable interest and we have gained a number of useful contacts – we consider it our most effective publication to date in reaching those who have the capacity to do something effective about this social problem.
Professor Paul Mullen, Monash University, Australia 'Cutting edge: Stalking', September 2000

I found it particularly beneficial writing for The Psychologist*. I was surprised by the amount of international correspondence and reprint requests it generated, from as far afield as New Zealand and South Africa.*
Dr Rory O'Connor, University of Strathclyde 'Suicidal behaviour', January 2001

It was a good idea to participate in the debate, especially considering its topic – hearing voices. The debate reached a much wider audience than specialist publications do, and so it afforded me a contact with individuals who hear voices whom I would not encounter otherwise.
Dr Ivan Leudar, University of Manchester 'Head to head: Is hearing voices a sign of mental illness?', May 2001

Writing the article helped me to crystallise certain ideas and to think about ways of presenting them to a wider audience than usual. The broad readership meant that I was approached by academics and other agencies I would not usually come into contact with, with the view of forging research links.
Dr Andrew Scholey, University of Northumbria 'Cutting edge: Fuel for thought', May 2001

See p.447 for contributors' information