

Anatomical Segregation of Component Processes in an Inductive Inference Task

Vinod Goel

York University and Institute of Neurology

Raymond J. Dolan

Institute of Neurology and Royal Free Hospital School of Medicine

Abstract

■ Inductive inference underlies much of human cognition. The essential component of induction is hypothesis selection based on some criterion of relevance. The purpose of this study was to determine the neural substrate of inductive inference, particularly hypothesis selection, using fMRI. Ten volunteers were shown stimuli consisting of novel animals under two task conditions, and asked to judge whether all the animals in the set were the same type of animal. In one condition, subjects were given a rule that specified the criteria for “same type of animal”. In the other condition, subjects had to infer the rule without

instruction. The two conditions were further factored into easy and difficult components. Rule inference was specifically associated with bilateral hippocampal activation while the task by difficulty interaction was associated with activation in right lateral orbital prefrontal cortex. We interpret the former in terms of semantic encoding of novel stimuli, and the latter in terms of hypothesis selection. Thus, we show an anatomical dissociation between task implementation and task difficulty that may correspond to a critical psychological distinction in the processes necessary for inductive inference. ■

INTRODUCTION







Human reasoning can be classified into two broad categories: deduction and induction. Deduction is an important but specialized form of reasoning which involves explicating information implicit in a closed system. Induction, on the other hand, is open-ended and takes one beyond the information contained in the premises. Induction is an ubiquitous, often effortless, process involved in many cognitive tasks, from perception, categorization, to explicit reasoning in problem-solving and decision-making.

Long a concern of philosophers, induction has been empirically studied by cognitive scientists for the past 25 years. Cognitive/computational models of induction typically view it as a form of hypothesis generation and selection, where one must search a large data base and determine which items of information are relevant and how they are to be mapped onto the present situation (Carbonell, Michalski, & Mitchell, 1983; Russell, 1986). The determination of hypothesis relevance is the crucial component of induction (Goodman, 1955), but little is known about underlying computational and neural mechanisms. For example, in reading about a recent discovery of dinosaur fossils in the tar pits of Alberta, we learn that the specimen in question, a Tyrannosaurus, had 8-in. long, razor-sharp teeth, and died by drowning in the tar pits. From this statement, we automatically infer

that all Tyrannosaurus had large, razor-sharp teeth. But we do not infer that all Tyrannosaurs died by drowning in tar pits. The logic of the inference, while invalid, is identical in both cases. However, the former inference is plausible; the latter is not. We recognize that “having large, sharp teeth” is a relevant property for generalization across members of a species, but that mode of death is a matter of individual accident. The puzzle of induction is, to a large extent, the question of how we make these judgments of relevance.

In the present study, we address the neuroanatomical correlates of inductive inference in an (novel) animal categorization paradigm. Ten subjects underwent fMRI (functional magnetic resonance imaging) scanning while shown sets of novel animals (called Caminalcules, after their creator, J. H. Camin), with the task requirement of determining if all of the animals in the set were of the same type. In one condition, subjects were given the following arbitrary rule for determining *same type of animal*: If the animals have the same tail and abdomen conditions (that is, the presence or absence of a tail and the same shaped abdomen), then they are the “same type of animal.” In a second condition (induction), subjects had to generate/infer their own rule for determining “same type of animal”. In addition to this manipulation, we also introduced another factor involving task difficulty (easy vs. hard). Task difficulty was

Figure 1. Sample stimuli from easy and hard conditions. Stimuli size was controlled in the actual presentation.

Tasks	Easy Condition Stimuli	Hard Condition Stimuli
Same type of animal?		
Same type of animal based on the given rule?		
Are the animals arranged in a straight line?		

manipulated by varying the degree of similarity between the animal figures. In the easy condition, the stimuli were either very different, or exact duplications while in the hard condition, there were subtle variations between the figures. Exemplar stimuli are presented in Figure 1. A low-level perceptual task, where subjects were asked to determine whether the set of animal stimuli were arranged in a straight line or a triangle, served as a common baseline. Subjects responded by pressing *yes* or *no* buttons on a response box.

Our cognitive analysis of the task is illustrated in Figure 2. It is assumed that visual processing and shape registration of the stimuli is obligatory. However, in the rule inference condition routine semantic encoding/

processing of the stimuli will fail (because the animals are all novel). Behavioural data indicate that, upon failure, subjects will fall back upon encoding readily recognizable features, such as horns on head, size and shape of abdomen, structure of arms and legs, body markings, and so on. (Sokal, 1974). They must subsequently determine which of these features are relevant for classification. Once the relevant features have been determined, subjects have a rule of classification that they can apply to the stimuli.

In the rule application condition, subjects are given a rule which makes reference to familiar, easily recognizable features such as abdomen and tails. They merely need to encode the familiar features and apply the given rule. If all the animals have the specified features, subjects respond “yes”, otherwise “no”. Subjects are not confronted with a novel classification task. Neither are they required to select a combination of features (rule) on which to base the classification. In fact, the optimal strategy in this condition is to suppress novel classification (and the semantic processing it entails) and focus on the familiar features specified by the rule.

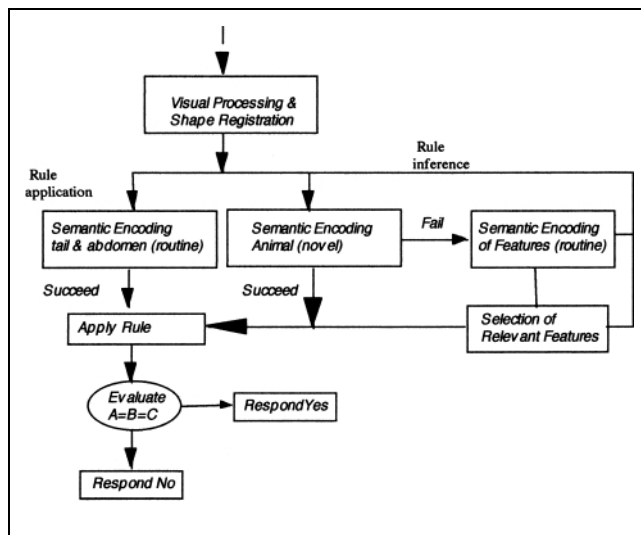


Figure 2. Analysis of task.

RESULTS

Subjects' responses and reaction times are reported in Table 1. Reaction-time data indicated a significant main effect of task difficulty. Subjects took a mean of 2122 ms (SD=399) to do the task in the easy condition compared to a mean of 2442 ms (SD=358) in the hard condition ($F(1, 9)=24.4, p<.05$). The main effect of the task was not significant (2297 ms (SD=365) for the induction task compared to 2266 ms (SD=456) for the rule application task). There was no significant interaction in the task by difficulty reaction times. Subjects'

Table 1. Subjects' Reaction Times and Scores (See Footnote 1)

	<i>Easy Condition</i>	<i>Hard Condition</i>
<i>Induction</i>		
Rts (ms) (SD)	2109 (407)	2485 (192)
Scores (%) (SD)	91.2 (11.6)	71.6 (17.7)
<i>Rule Application</i>		
Rts (ms) (SD)	2134 (413)	2399 (480)
Scores (%) (SD)	86.8 (17.6)	70.8 (9.9)
<i>Perceptual Baseline</i>		
Rts (ms) (SD)	1635 (378)	1634 (421)
Scores (%) (SD)	96.82 (5.3)	98.0 (3.9)

response profiles were very similar.¹ There was a significant main effect of task difficulty with subjects' responding correctly on 89.0% (SD=14.7) of easy trials and 71.2% (SD=14.0) of hard trials ($F(1, 9)=21.9$, $p<.05$). The main effect of the task was not significant, with subjects scoring 81.4% (SD=17.7) on induction trials, compared to 78.8% (SD=16.2) on rule application trials. Neither was there a task by difficulty interaction in the scores. The scores and reaction times for the individual cells of the factorial design are shown in Table 1. The perceptual baseline was designed to be independent of the difficulty manipulation, and neither mean reaction times (1635 ms (SD=378) with easy stimuli and 1634 ms (SD=421) with hard stimuli), nor mean task scores (93.2% (SD=5.3) with easy stimuli and 93.2% (SD=5.0) with hard stimuli) were affected by the manipulation.

The fMRI results are summarized in Table 2. The simple effect of induction [induction–perceptual baseline] was associated with bilateral activation in the cerebellum, occipital cortex (BA 18, BA 19, BA 37), superior parietal lobe (BA 7), and prefrontal cortex (BA 44, 46) (Figure 3). However, there was a strong hemispheric asymmetry (R>L) for the parietal and prefrontal cortex. The simple effect of rule application condition [rule application–perceptual baseline] activated a similar network, but without the hemispheric asymmetry and with the addition of the medial prefrontal cortex (BA 8) (Figure 3).

The main effect of task (induction–rule application) revealed significant activation in bilateral hippocampus (L>R) (see Figure 4). Analysis of the response profile indicated that this effect is largely attributable to a deactivation, relative to perceptual baseline, in the rule application condition. However, the inference condition was associated with enhanced activation in the left hippocampus, compared to the perceptual baseline, an

effect not seen for the right hippocampus. The reverse comparison (rule application–induction) revealed significant activation in presupplementary motor area (BA 8) (see Figure 5).

The main effect of task difficulty did not reveal any significant regions of activation (after correction for multiple comparison), though an activation in the cingulate (BA 4, BA 30, BA 46) did approach significance ($Z=3.68$) in the difficult-minus-easy comparison.

The task-by-difficulty interaction ([hard induction–hard rule application]–[easy induction–easy rule application]) was associated with a significant effect in the right inferior prefrontal cortex (BA 47, BA 11) (see Figure 6). Estimation of the condition-specific response profile of the right inferior prefrontal cortex indicated that enhanced activation in the hard inference condition, along with decreased activation in the rule application condition, contributed to the effect. The reverse interaction ([easy induction–easy rule application]–[hard induction–hard rule application]) did not reveal any regions of significant activation.

DISCUSSION

The behavioral results confirm that subjects engaged in the task in the expected manner and verified the effectiveness of the difficulty manipulation.

The simple effects of rule induction and rule application revealed a similar network, but with a different hemispheric profile. The rule induction condition showed a predominant pattern of RH activation, while the profile of activation for rule application was more bilateral. The greater LH activation in the rule application condition (as compared to rule induction), may reflect the requirement of maintaining the rule on-line, while carrying out feature-based searches. The involvement of the right inferior and dorsolateral prefrontal cortex (I/DLPFC) in the induction condition is particularly interesting in light of previous studies of deductive reasoning, problem-solving, and language, and is discussed below.

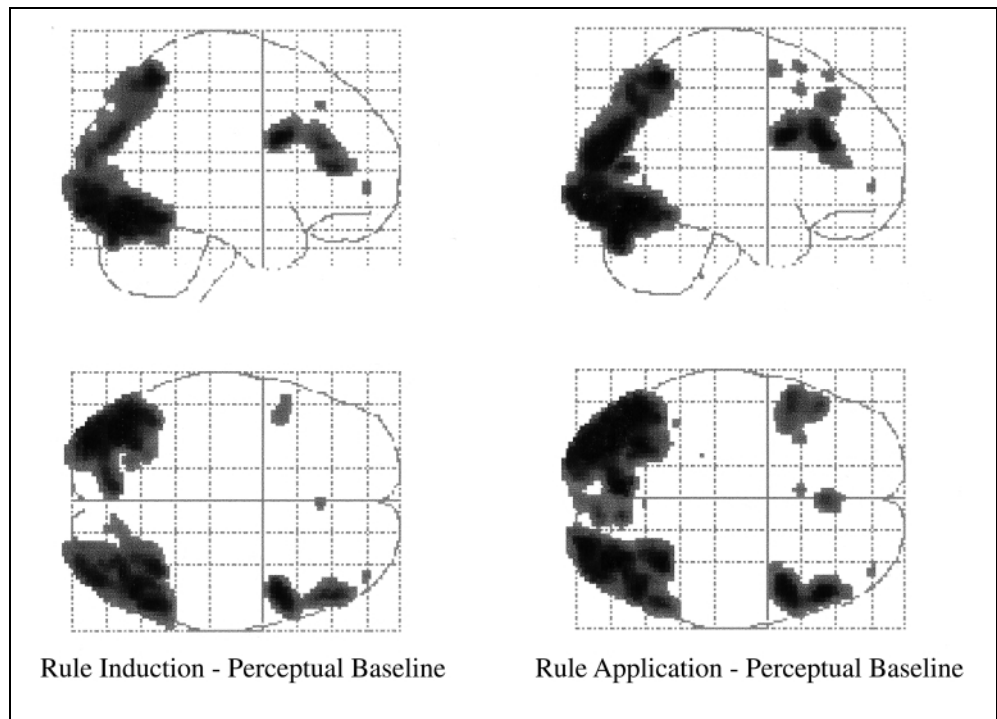
Our results for the simple effects of rule application overlap with the results of a recent study of rule-based and similarity-based categorization by Smith, Patalano, and Jonides (1998). Their rule-based condition was very close to our rule application condition, so not surprisingly, there is considerable overlap in activation patterns (including BA 18, BA 19, BA 7, BA 46) for this condition. Their “similarity” condition was an exemplar-based memory task, very different from our induction condition. This may account for the differences in the other comparisons.

The main effect of the rule inference minus the rule application condition was associated with bilateral hippocampal activation. Before considering the likely functional significance of this finding, it is worthwhile to consider the relevant psychological processes. It is

Table 2. Location and Characteristics of the Brain Regions that Remained Significantly Active After Individual Comparisons

<i>Location (Brodmann's Area)</i>	<i>Size (k)</i>	<i>MNI Coordinates</i>			<i>Z-score</i>
		<i>X</i>	<i>Y</i>	<i>Z</i>	
<i>Induction–Perceptual Baseline</i>					
Cerebellum	3583	–38	–76	–16	7.41
Lt. fusiform gyrus (BA 18)		–32	–90	–12	7.09
Lt. occipital gyrus (BA 19)		–44	–70	–10	6.98
Rt. superior parietal lobule (BA 7)	4717	34	–56	58	7.00
Rt. occipital gyrus (BA 19)		48	–66	–16	6.89
Rt. inferior temporal gyrus (BA 37)		54	–58	–14	6.86
Rt. inferior frontal gyrus (BA 44)	894	48	10	26	6.93
Rt. inferior frontal gyrus (BA 46)		50	38	12	6.29
Rt. middle frontal gyrus (BA 46)		50	32	22	6.15
Lt. superior parietal lobule (BA 7)	237	–24	–62	46	5.76
Lt. superior parietal lobule (BA 7)		–22	–68	58	5.53
Lt. inferior frontal gyrus (BA 44)	81	–44	10	28	5.18
<i>Rule Application–Perceptual Baseline</i>					
Cerebellum	6789	–28	–92	–12	7.44
Lt. fusiform gyrus (BA 18)		–38	–76	–16	7.42
Lt. middle occipital gyrus (BA 18)		–30	–84	16	7.36
Rt. middle occipital gyrus (BA 18)	5043	32	–84	22	7.20
Lt. superior parietal lobule (BA 7)		34	–56	56	7.11
Rt. fusiform gyrus (BA 18)		24	–96	–10	7.04
Rt. middle frontal gyrus (BA 46)	1083	50	30	24	7.04
Rt. inferior frontal gyrus (BA 44, 9)		48	10	26	7.02
Rt. superior frontal sulcus (BA 6)	31	36	4	60	5.06
Lt. inferior frontal gyrus (BA 9)	727	–52	12	32	6.07
Lt. middle frontal gyrus (BA 46)		–48	28	24	6.00
Medial prefrontal cortex (BA 8)	162	0	30	44	5.87
<i>Induction–Rule Application</i>					
Lt hippocampus	245	–20	–14	–16	4.55
Rt hippocampus	33	26	–18	–20	3.29
<i>Rule Application–Induction</i>					
Lt presupplementary motor area (BA 8)	405	–4	34	58	5.16
<i>(Hard Induction–Easy Induction)–(Hard Rule Application–Easy Rule Application)</i>					
Rt. inferior frontal gyrus (BA 47, BA 11)	338	34	28	–2	4.71
		52	42	–18	4.17
		38	40	–14	3.92

Figure 3. Simple effects of rule induction [induction–perceptual baseline] and rule application conditions [rule application–perceptual baseline].



generally accepted that inductive inferences are a function of semantic content, rather than formal features of stimuli (Rescher, 1980). Therefore, stimuli must be encoded at the semantic level, a process that requires a deeper level of processing than form-based encoding (Craik & Lockhart, 1972). Furthermore, there is a strong novelty component introduced by the nature of the task. In the rule inference condition, subjects are confronted with novel animal stimuli that do not conform to any a priori classification scheme, and are forced to attempt semantic encoding (as a prerequisite to inference).

In light of these considerations, one explanation for the left hippocampal activation is the enhanced semantic encoding of novel stimuli that is a prerequisite for inductive inference. Recent neuroimaging studies indicate that an hippocampal response during memory encoding is modulated by stimulus novelty (Dolan & Fletcher, 1997; Schacter & Wagner, 1999; Stern et al., 1996). What constitutes novelty is, to a large extent, relative to context. We would argue that in our experimental design, the stimuli have greater contextual novelty in the inference condition because, as subjects try to encode the figures *as animals*, they are unable to incorporate them into a pre-existing classification scheme. Furthermore, examination of the parameter estimates associated with hippocampal activation (Figure 5) indicates that a deactivation of the hippocampus (compared to perceptual baseline) during the rule application condition contributes to the main effect under consideration. A possible explana-

tion for this is that in the rule application condition, there is an absence of an obligatory attempt to encode the novel stimuli (as animals) in favour of attending to familiar *features* of the stimuli (for example, tails and abdomen) specified in the rule application condition.

The reverse subtraction (rule application–rule inference) revealed activation of the presupplementary motor area (BA 8), a region implicated in anticipation/preplanning of motor activity. One explanation of this activation is that since the subjects know the rule prior to stimulus presentation (and the rule is the same for all conditions), they anticipate and prepare a response to features specified in the rule. In the rule inference condition, this anticipation is not possible, because the rule is unknown.

In addition to semantic encoding, induction also requires that subjects access their (implicit) knowledge of categorizing animals. Most subjects categorize by picking out animal features, such as shape of abdomen, number of appendages, structure of appendages, body markings, positioning of eyes, and so on (Sokal, 1974) (that is, hypothesis generation), and then determining the relevance of each feature (that is, hypothesis selection). Such selection is based upon personal experience and world knowledge. For example, having spots on the abdomen may be a reasonable species-differentiating property. However, dalmatians have spots but poodles do not, yet they are both dogs. Consequently, subjects must decide whether a given property is relevant in a particular instance or not. Recall that in the rule applica-

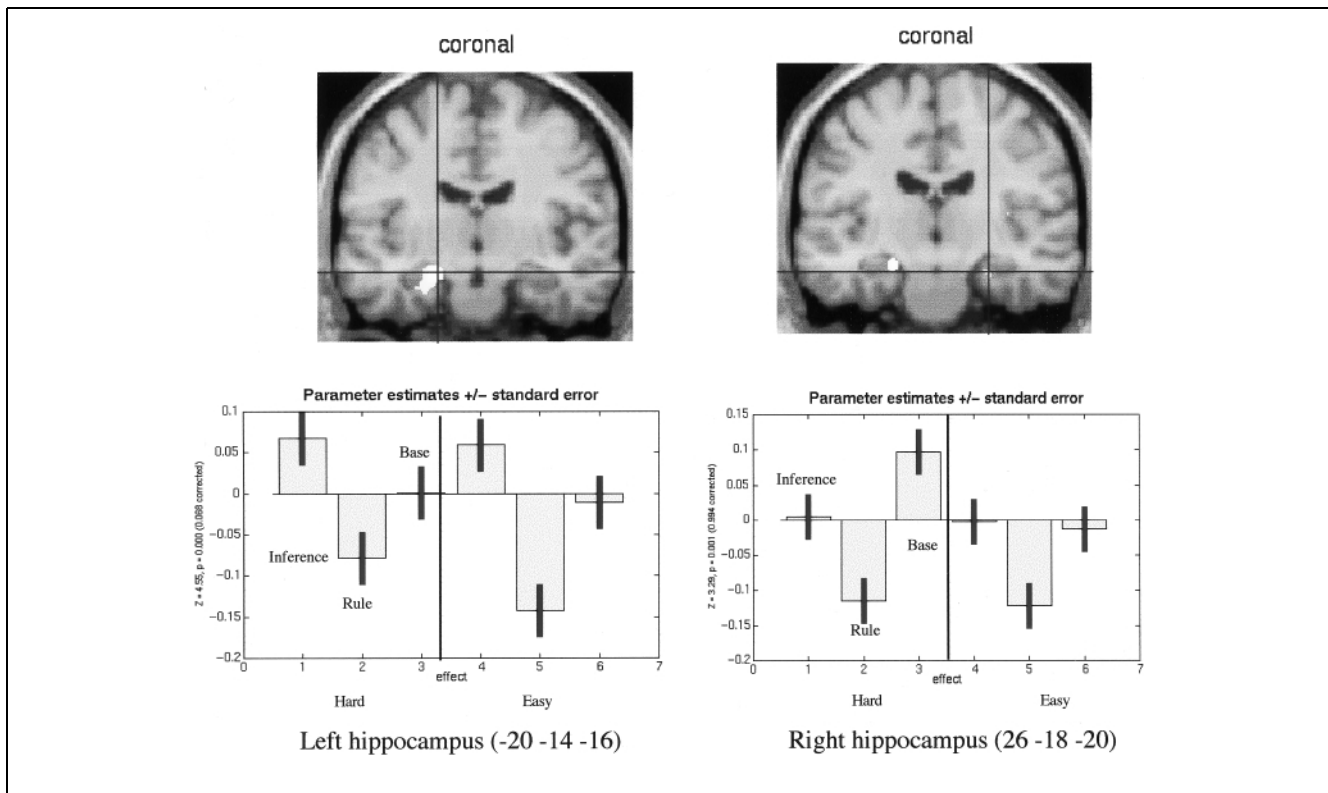


Figure 4. Main effect of rule induction (rule induction–rule application).

tion condition, subjects are told which features are relevant, and they simply have to determine whether the features are present in all three of the animals. They need not make a decision about the relevance of the features.

An analysis of the task by difficulty interaction revealed that increased difficulty in the induction condition activated the right lateral orbital prefrontal cortex (BA 47, BA 11).² While induction requires both generation and selection of properties, we favour a hypothesis selection account of this activation. Our rationale for implicating only selection is based on task demands and anatomical considerations. Generation of features is a relatively straightforward matter in our task because all the features are present in the stimuli. The challenging issue is to determine the relevance of each feature (that is, hypothesis selection). As we move from the easy induction condition to the hard induction condition, the number or presence of features does not change. However, there is an increase in difficulty in evaluating the hypotheses. In the easy condition, the animals are either identical or very different (see Figure 1). In the hard condition, there are subtle variations in the features that make the evaluation/selection of the hypotheses more difficult (see Figure 1 and Table 1). In keeping with the above, the anatomical regions activated are not those previously associated with memory retrieval but they do overlap with regions reported by Elliott and Dolan (1998) in a hypothesis selection task that had no

right or wrong answers. We, therefore, attribute the right prefrontal activation to hypothesis selection rather than to hypothesis generation.

Our right prefrontal activation in the induction condition is particularly interesting in light of the fact that imaging and lesion studies of *deductive* reasoning have consistently implicated the *left* I/DLPFC (Gazzaniga, 1985; Gazzaniga & Smylie, 1984; Goel, Gold, Kapur, & Houle, 1997a; Goel, Gold, Kapur, & Houle, 1998; Read, 1981). Deduction differs from induction in that it is a well-defined, closed-ended task with definite right or wrong answers, while induction is an ill-defined, open-ended task with no right or wrong answers. In conjunction with these studies, our results suggest that the right hemisphere may have a special role to play in open-ended inference tasks with no right or wrong answers. Consistent with our interpretation, a number of lesion studies have shown that RH lesions to the prefrontal cortex can selectively impair patient ability to engage in open-ended tasks, while leaving conventional neuropsychological measures intact (Eslinger & Damasio, 1985; Goel & Grafman, in press; Goel, Grafman, Tajik, Gana, & Danto, 1997b; Penfield & Evans, 1935). Our results are also broadly consistent with a series of language studies that demonstrate that while much of syntactic and semantic processing involves the LH, the RH has a special role to play in analogical and metaphorical transfer involving inductive inference (Bottini et al., 1994; Brownell,

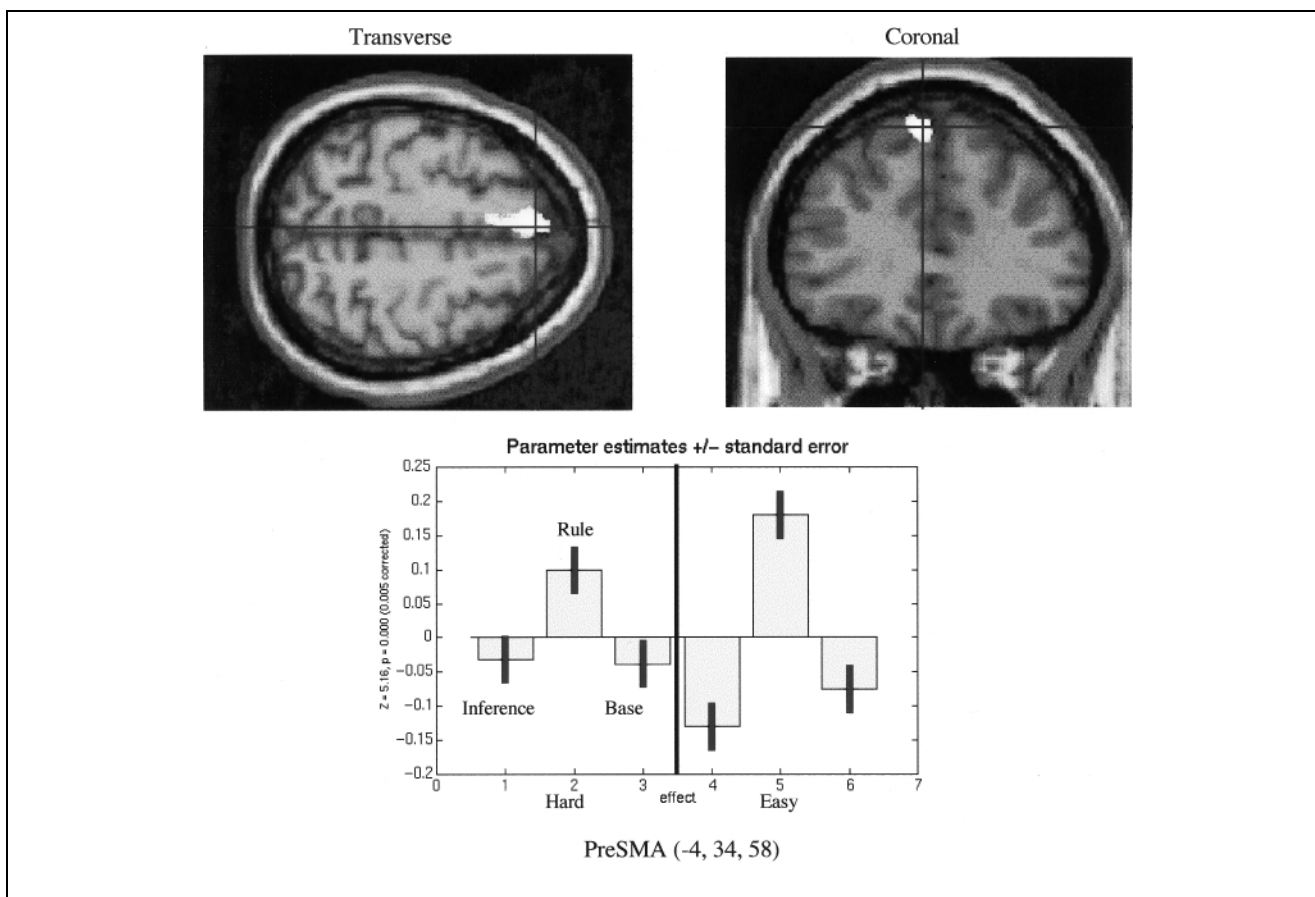


Figure 5. Main effect of rule application (rule application—rule induction).

Simpson, Bihrlé, Potter, & Gardner, 1990; Wapner, Hamby, & Gardner, 1981).

In conclusion, we have demonstrated two distinct patterns of activation associated with an inductive inference task. Task engagement led to hippocampal activation, which we interpret in terms of semantic encoding of novel stimuli. We attribute the right lateral orbital prefrontal cortex activation to determination of hypothesis relevance in an open-ended task with no right or wrong answers. The former is unaffected by task difficulty, whereas determination of hypothesis relevance is clearly affected and is associated with a significant interaction. Thus, the psychological distinction between components of induction is underwritten by an anatomical dissociation.

METHOD

Subjects

Ten right-handed normal subjects (seven males and three females), with a mean age of 26.4 years and mean education level of 15.9 years, volunteered to participate in the study. All subjects gave informed consent, and the study was approved by the Joint National Hospital for Neurology and Neurosurgery/Institute of Neurology Ethics Committee.

Stimuli

A total of 14 black and white line drawings of novel animals were adopted from Sokal (1974). These are all fictitious animals but have superficial similarities to real animals. Subjects were told (and accepted) that they were real animals. Additional stimuli were generated from the set of 14 by either (i) duplication, (ii) spatial transformations such as elongation and warping, (iii) deleting or adding body parts such as tails and appendages, and (iv) changing body markings (for example, from spots to stripes). Examples of these transformed animals appear in Figure 1. Each stimulus array presentation consisted of three animals arranged in either a straight line or triangle. A total of 60 different screens were presented.

Experimental Design

A 2×2 factorial design was used. The first factor involved either inferring and applying a categorization rule or just applying a given categorization rule. In the “rule inference”³ or “induction” condition, subjects were asked whether “all three animals are of the same type”. They were told to use their general knowledge of categorizing familiar animals such as cats and dogs to do the task. In the “rule application condition”

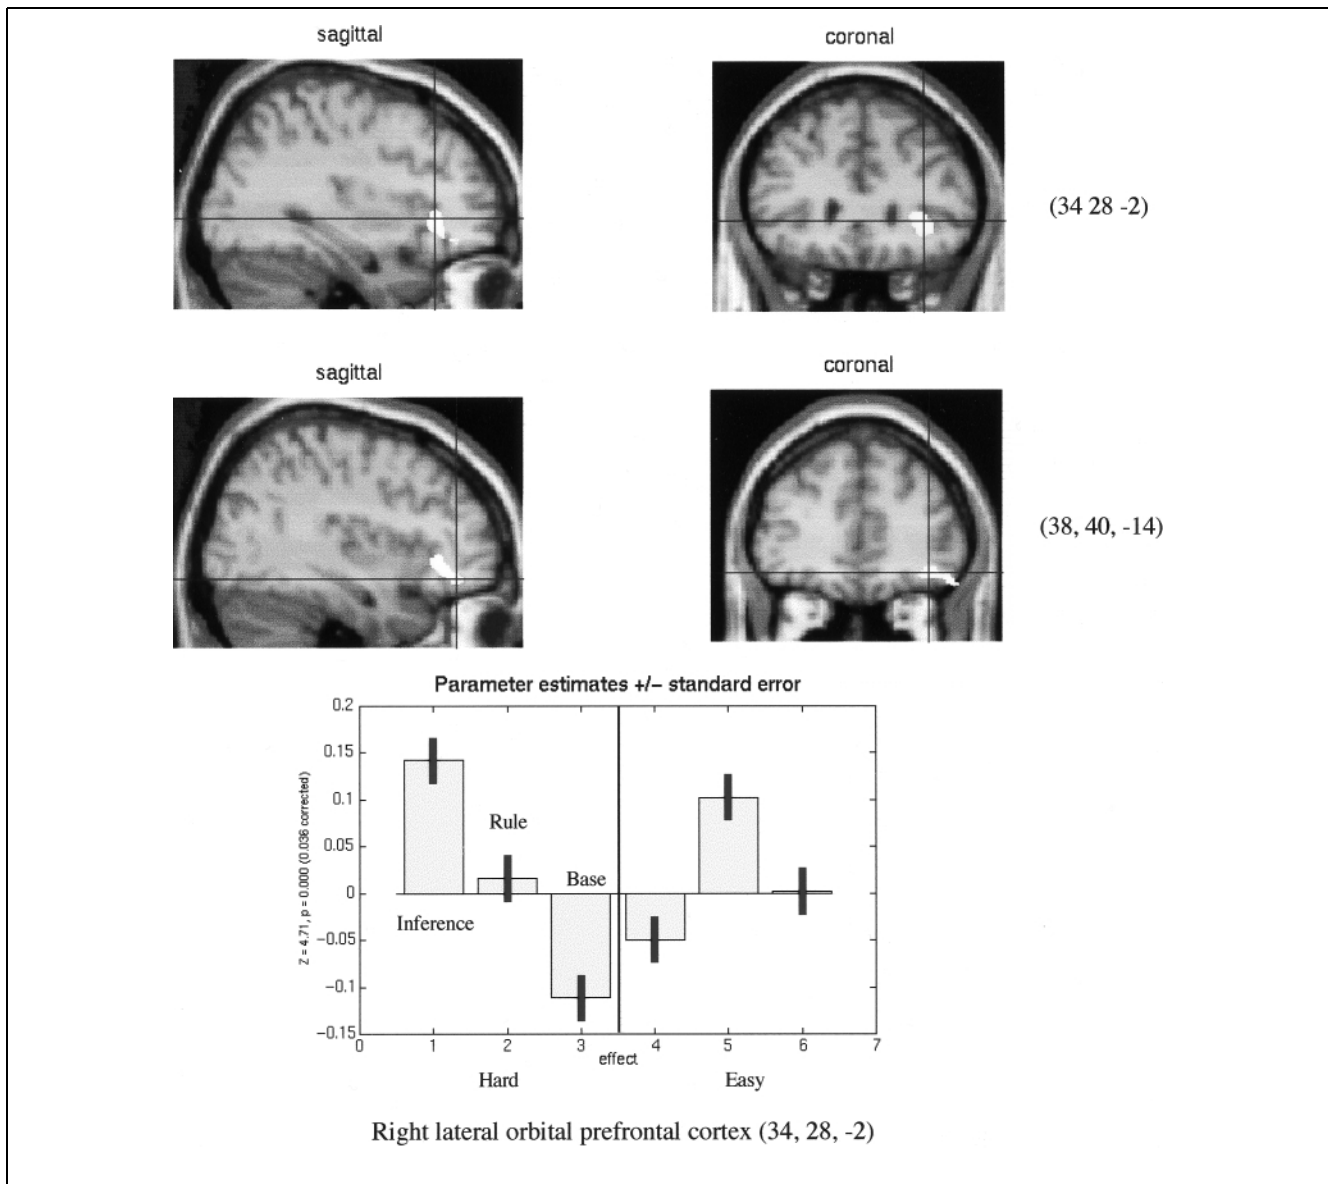


Figure 6. Task-by-difficulty interaction ([hard rule induction–hard rule application]–[easy rule induction–easy rule application]).

subjects were given the following rule to apply: “If the animals have the same tail and abdomen conditions (that is, the presence or absence of a tail and the same shaped abdomen), then they are the same type of animal”. The second factor, task difficulty, was manipulated by varying the degree of similarity between the figures. In the easy condition, the stimuli were either very different, or exact duplications while in the hard condition, there were more subtle variations between the figures. Exemplar stimuli from the easy and hard conditions are presented in Figure 1 with each level consisting of 30 different stimuli screens. A perceptual baseline condition was also utilized in which subjects were asked whether the three animals were arranged in a straight line or a triangle.

Half of the items in each condition required a “yes” response; half required a “no” response. Subjects were

asked to respond as quickly as possible and instructed to move onto the next trial if the stimuli advanced before they could respond. They responded by pressing one of two keys on a keyboard. Subjects reviewed example stimuli from each condition prior to being scanned to ensure that they understood the task. Subjects’ compliance with study instructions was determined through debriefing and reaction time scores.

fMRI Scanning Technique

A 2T Siemens VISION system (Siemens, Erlangen, Germany) was used to acquire T1 anatomical volume images (1×1×1.5 mm voxels) and 48 T2*-weighted echoplanar images (64×64 3×3 mm pixels, TE=40 ms) with blood oxygenation level dependent (BOLD) contrast. Echoplanar images 1.8-mm thick were acquired axially every 3 mm, positioned to cover the whole brain. Data were

recorded in one session. A total of 306 volume images were acquired continuously with a repetition time (TR) of 4.1 sec/vol. The first six volumes were discarded to allow for T1 equilibration effects.

The six conditions were blocked in a box-car design and presented in the following order: hard rule inference and application, hard rule application only, perceptual baseline, easy rule inference and application, easy rule application only, and perceptual baseline. Each condition corresponded to an epoch consisting of five trials, presented for 5000 ms with an interstimulus interval of 740 ms. Trials were preceded by a rest/instruction screen of 12,300 ms duration. This gave a total epoch duration of 41 sec, during which time 10 image volumes were acquired (seven during task conditions; three during rest/instruction screens). The scanner was synchronized with the presentation of the first trial in each epoch. The epoch cycle was repeated six times for a total of 30 trials per condition. The ordering of the conditions was fixed, but the distribution of stimuli across conditions was randomized to control for novelty and multiple exposure effects.

Data Analysis

Data were analyzed using Statistical Parametric Mapping (SPM 97d) (Friston et al., 1995). All volumes were spatially realigned to the first volume (head movement was <2 mm in all cases) and temporally realigned to the AC-PC slice, to account for different sampling times of different slices. A mean image created from the realigned volumes was co-registered with the structural T1 volume and the structural volumes spatially normalized to the Montreal Neurological Institute brain template (Evans et al., 1993) using nonlinear basis functions (Ashburner & Friston, 1999). The derived spatial transformation was then applied to the realigned T2* volumes, which were finally spatially smoothed with a 10 mm FWHM isotropic Gaussian kernel (in order to make comparisons across subjects and to permit application of random field theory for corrected statistical inference (Worsley & Friston, 1995)). The resulting time series across each voxel were high-pass filtered with a cut-off of 120 sec, using cosine functions to remove section-specific low frequency drifts in the BOLD signal. Global means were normalized by proportional scaling to a Grand Mean of 100 to remove confounding effects of whole brain activity,⁴ and the time series temporally smoothed with a 4 sec FWHM Gaussian kernel to swamp small temporal autocorrelations with a known filter.

Condition effects at each voxel were estimated according to the general linear model and regionally specific effects compared using linear contrasts. Each contrast produced a statistical parametric map of the *t*-statistic for each voxel, which was subsequently transformed to be

unit normal *Z*-distribution. The activations reported survived both a voxel-level and cluster-level correction of $p < .05$ ($Z > 4.60$) using a random effect model. An exception has been made in the case of anatomical symmetry. Where an anatomical structure is significantly active in one hemisphere, we report any activation in the corresponding structure in the other hemisphere, even if it does not reach threshold.

Acknowledgments

We are indebted to Chris Frith and Karl Friston for advice in the design and analysis phases (respectively) of the study and to Richard Frackowiak for internal review of this manuscript. We also thank two anonymous reviewers of this journal for considered, detailed comments that served to improve the manuscript. VG was supported in part by a York University Faculty of Arts Fellowship and NSERC and SSHRC grants. RJD is supported by the Wellcome Trust.

Reprint request should be sent to Vinod Goel, Department of Psychology, University of Aberdeen, Kings College, Old Aberdeen, Aberdeen, AB24 2UB, UK or R.J. Dolan at the Wellcome Department of Cognitive Neurology, Institute of Neurology, Queens Square, London WC1N 3BG, UK.

Notes

1. As there were no definitive correct responses in the induction task condition, subjects' responses for induction trials were compared to the responses of a single subject who did the task outside the scanner without time constraints.
2. Interestingly, this region does not achieve significant levels of activation in the main effect of task. This may reflect the fact that the easy rule inference condition was extremely easy, with subjects performing at 91% accuracy (see Table 1).
3. Strictly speaking it is a "rule inference and application" condition. However, to avoid confusion with the "rule application" condition, we refer to it as simply a "rule inference condition."
4. This means that our activations are regionally specific responses relative to whole brain activity. Normalization of the data by the global signal may attenuate differences between control and signals (Aguirre, Zarahn, & D'Esposito, 1998). To minimize this concern we have: (i) acquired whole brain data; (ii) reported only activations that survived both voxel-level and cluster-level correction of $p < .05$ ($Z > 4.60$) using a random effect model; and (iii) where activations are driven by relative deactivations in the relevant control condition, this is duly acknowledged.

REFERENCES

- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). The inferential impact of global signal covariates in functional neuroimaging analyses. *Neuroimage*, 8(3), 302–306.
- Ashburner, J., & Friston, K. J. (1999). Nonlinear spatial normalization using basis functions. *Human Brain Mapping*, 7(4), 254–266.
- Bottini, G., Corcoran, R., Sterzi, R., Paulesu, E., Schenone, P., Scarpa, P., Frackowiak, R. S. J., & Frith, C. D. (1994). The role of the right hemisphere in the interpretation of figurative aspects of language: A positron emission tomography activation study. *Brain*, 117, 1241–1253.

- Brownell, H. H., Simpson, T. L., Bihrlé, A. M., Potter, H. H., & Gardner, H. (1990). Appreciation of metaphoric alternative word meanings by left and right brain-damaged patients. *Neuropsychologia*, *28*, 375–383.
- Carbonell, J. G., Michalski, R. S., & Mitchell, T. M. (Eds.). (1983). *Machine learning: An artificial intelligence approach*. Palo Alto, California: Tioga Press.
- Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, *11*, 671–684.
- Dolan, R. J., & Fletcher, P. C. (1997). Dissociating prefrontal and hippocampal function in episodic memory encoding. *Nature*, *388*, 582–585.
- Elliott, R., & Dolan, R. J. (1998). Activation of different anterior cingulate foci in association with hypothesis testing and response selection. *Neuroimage*, *8*, 17–29.
- Eslinger, P. J., & Damasio, A. R. (1985). Severe disturbance of higher cognition after frontal lobe ablation: Patient EVR. *Neurology*, *35*, 1731–1741.
- Evans, A. C., Collins, D. L., Mills, S. R., Brown, E. D., Kelly, R. L., & Peters, T. M. (1993). 3D statistical neuroanatomical models from 305 MRI volumes. *Proceedings Institute of Electrical and Electronics Engineers—Nuclear Science Symposium and Medical Imaging Conference*, 1813–1817.
- Friston, K., Holmes, A., Worsley, K., Poline, J. -B., Frith, C., & Frackowiak, R. (1995). Statistical parametric maps in functional imaging: A general approach. *Human Brain Mapping*, *2*, 189–210.
- Gazzaniga, M. S. (1985). *The Social brain*. NY: Basic Books.
- Gazzaniga, M. S., & Smylie, C. S. (1984). Dissociation of language and cognition. *Brain*, *107*, 145–153.
- Goel, V., & Grafman, J. (in press). The role of the right prefrontal cortex in ill-structured problem solving. *Journal of Cognitive Neuropsychology*.
- Goel, V., Gold, B., Kapur, S., & Houle, S. (1997a). The seats of reason: A localization study of deductive and inductive reasoning using PET (O15) blood flow technique. *NeuroReport*, *8*(5), 1305–1310.
- Goel, V., Grafman, J., Tajik, J., Gana, S., & Danto, D. (1997b). A study of the performance of patients with frontal lobe lesions in a financial planning task. *Brain*, *120*, 1805–1822.
- Goel, V., Gold, B., Kapur, S., & Houle, S. (1998). Neuroanatomical correlates of human reasoning. *Journal of Cognitive Neuroscience*, *10*(3), 293–302.
- Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.
- Penfield, W., & Evans, J. (1935). The frontal lobe in man: A clinical study of maximum removals. *Brain*, *58*, 115–133.
- Read, D. E. (1981). Solving deductive-reasoning problems after unilateral temporal lobectomy. *Brain and Language*, *12*, 116–127.
- Rescher, N. (1980). *Induction: An essay on the justification of inductive reasoning*. Pittsburgh: University of Pittsburgh Press.
- Russell, S. J. (1986). Preliminary steps toward the automation of induction. *Proceedings of the AAAI*, 119–126.
- Schacter, D. L., & Wagner, A. D. (1999). Medial temporal lobe activations in FMRI and PET studies of episodic encoding and retrieval. *Hippocampus*, *9*(1), 7–24.
- Smith, E. E., Patalano, A. L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition*, *65*, 167–196.
- Sokal, R. R. (1974). Classification: Purposes, principles, progress, prospects. *Science*, *185*, 1115–1123.
- Stern, C. E., Corkin, S. C., Gonzalez, R. G., Guimaraes, A. R., Baker, J. R., Jennings, P. J., Carr, C. A., Sugiura, R. M., Vedantham, V., & Rosen, B. R. (1996). The hippocampal formation participates in novel picture encoding: Evidence from functional magnetic resonance imaging. *Proceedings National Academy of Science, USA*, *93*, 8660–8665.
- Wapner, W., Hamby, S., & Gardner, H. (1981). The role of the right hemisphere in the apprehension of complex linguistic materials. *Brain and Language*, *14*, 15–33.
- Worsley, K. J., & Friston, K. J. (1995). Analysis of fMRI time-series revisited-again. *Neuroimage*, *2*, 173–181.