Cognition, 50 (1994) 347–362
0010-0277/94/$07.00

Levels of causal understanding in chimpanzees and children

David Premack*, Ann James Premack
Laboratoire de Psycho-Biologie du Développement, CNRS, 41 rue Gay Lussac, 75005 Paris, France

Abstract

We compare three levels of causal understanding in chimpanzees and children: (1) causal reasoning, (2) labelling the components (actor, object, and instrument) of a causal sequence, and (3) choosing the correct alternative for an incomplete representation of a causal sequence. We present two tests of causal reasoning, the first requiring chimpanzees to read and use as evidence the emotional state of a conspecific. Despite registering the emotion, they failed to use it as evidence. The second test, comparing children and chimpanzees, required them to infer the location of food eaten by a trainer. Children and, to a lesser extent, chimpanzees succeeded. When given information showing the inference to be unsound – physically impossible – 4-year-old children abandoned the inference but younger children and chimpanzees did not. Children and chimpanzees are both capable of labelling causal sequences and completing incomplete representations of them. The chimpanzee Sarah labelled the components of a causal sequence, and completed incomplete representations of actions involving multiple transformations. We conclude the article with a general discussion of the concept of cause, suggesting that the concept evolved far earlier in the psychological domain than in the physical.

*Corresponding author.
The data reported here were collected at the University of California, Santa Barbara, and the University of Pennsylvania Primate Facility. We are greatly indebted to Guy Woodruff, who participated in all phases of the research and would be a co-author if we knew his whereabouts and could obtain his permission. We are also indebted to the many students, graduate and undergraduate, at both institutions who assisted in the care and testing of the chimpanzees.

SSDI 0010-0277(94)00617-T

Introduction

In this paper we compare three levels of causal understanding in chimpanzees and children.

At the deepest level, the individual engages in causal reasoning, solving problems in which he sees the outcome of a process but not its cause, and must infer or reconstruct the missing events. At an intermediate level the individual analyses intact causal sequences into their components and labels them. He must label the actor, object, and instrument of the action. The ability to carry out this task is a prerequisite for causal reasoning; if one cannot identify the separate parts of an intact sequence, one cannot identify the missing part of an incomplete sequence.

At the most superficial level, the individual must complete an incomplete representation of a causal action, by selecting the correct alternative. Since the alternatives are all visible this task is the least demanding.

Chimpanzees have been shown to be capable of all but the deepest level (Premack 1976, 1983), though they have been shown capable of analogical reasoning. Not only do they complete incomplete analogies and make same/different judgments about exemplars that are and are not analogies (Gillan, Premack, & Woodruff, 1981), they also construct analogies from scratch (Oden & Premack, unpublished data). There is evidence (Gillan, 1981), albeit inconclusive, that they can do transitive inference. But there is little indication that they are capable of "Sherlock Holmes" type reasoning, that is, causal reasoning.

In causal reasoning, an outcome – a corpse on the floor – is presented, but the cause is not. A human confronted with this scene would ask immediately, "Who did it? How? When, where and why?" He would answer these questions by making inferences from the "evidence", which has two main sources: existing knowledge concerning the "corpse" and its past, and observations concerning the immediate circumstances. As an astute observer, Sherlock Holmes was a good reasoner because he had an uncanny sense of what was "relevant", detecting implications in what others dismissed as neutral facts.

Causal reasoning

In this paper we present two tests of causal reasoning: one conducted with a group of chimpanzees, another in which we compare chimpanzees and young children.

Subjects

The chimpanzees (*Pan troglodytes*) were African born; four were $3\frac{1}{2}$–$4\frac{1}{2}$ years old and Sarah was about 10 years old at the time of the study. Animals entered the laboratory as infants, were diapered and bottle-fed, trained extensively in match-to-sample, and (some) were taught an artificial language. They were maintained rather like middle-class children on three meals a day, snacks, and continuous access to water. When tested, some were "rewarded" with a preferred food while others were simply praised. Participating children came from Philadelphia schools, and varied in age from 3.8 to 4.5 years, with an average age of 4.1 years.

Reading emotional evidence

Four juvenile chimpanzees were tested in a simple two-step reasoning problem. We first trained them to run down a path to a consistently positive hidden object, then introduced them to occasional negative trials; that is, a rubber snake was substituted for the food on 15% of trials on a random schedule. The unpredictable negative trials profoundly changed the character of the chimpanzees' run. They no longer dashed full speed to the goal, but slowed midway, approaching the goal hesitantly.

We next offered the animals an opportunity to play Holmes, to escape the uncertainty of the negative trial by using the emotional state of an informant to infer what object it had encountered on its run. Now, before starting a run, each animal was placed in a holding room with an informant that had just completed a run. The informant, having encountered either food or the rubber snake on its run, was in either a positive or negative emotional state, and this state, we have reason to believe, was successfully communicated to the recipient.

Uninformed human judges shown videotapes of the informant could discriminate the state of the informant (ca. 98% correct). Beyond that, they could discriminate the recipient's state following its contact with the informant (ca. 70% correct). However, the "informed" chimpanzees seemed not to profit from this contact, for they accepted all opportunities to run, and did so in the same way whether: (1) the informant was in a positive state, (2) a negative state, or even when on control trials (3) had had no contact with an informant at all.

The holding room was adjacent to the runway. Animals were taken there prior to, and immediately after, a run (to serve as an informant). Every chimpanzee played both roles, that of informant and recipient, and had the opportunity to observe that its conspecifics too played both roles. The use of four animals permitted 12 possible recipient–informant pairs, all of which were used.

Under these conditions, the animals should have been able to infer that an informant's emotional state was the result of what it had found at the end of a run. No chimpanzee ever encountered another in the holding room whose emotional state was not owed to a run. Nonetheless, it still could not use the emotional state (which it registered at some level) as evidence from which to infer what the informant had encountered on its run.

Could the chimpanzee have made this inference but not have assumed that what the informant found was not a good prediction of what *it* would find – in other words, assumed it might be snakes for you but food for me? We cannot rule out this possibility, but a human in this circumstance would certainly explore the hypothesis that snakes for you means snakes for me as well.

Before testing the chimpanzees, we had assumed this was a simple problem, one that could be easily solved, and therefore could be used as a base condition on which to impose variations that would permit our answering fundamental questions about reasoning. Perhaps at $3\frac{1}{2}$–$4\frac{1}{2}$ years the chimpanzees were too young and could have solved the problem when older. The age at which children can solve this problem is not known, for one cannot test children with frightening objects.

Using location as evidence

In the next experiment, we tested both chimpanzees (the same group of four used in the previous problem), and two groups of children (10 in each group). The apes were tested in their outdoor compound, and the children in a classroom.

For the chimpanzees, we placed two opaque containers about 30 feet apart. These formed the base of a triangle with the chimpanzee at its apex, midway between the containers and 30 feet from the base. The chimpanzee was accompanied by a trainer. As the two watched, a second trainer placed an apple in one container and a banana in the other. Following this action, the accompanying trainer placed a blind in front of the chimpanzee so that the containers were no longer visible to it. The trainer distracted the animal for approximately 2 min before removing the blind.

What the subject now saw was the second trainer standing midway between the containers eating either an apple or a banana. Having eaten the fruit, the trainer left, and the chimpanzee was released. Each animal was given 10 trials, with an intertrial interval of about an hour. The two fruit were placed equally often in both containers, and the fruit eaten by the experimenter was counterbalanced over trials. We used apples and bananas because apes are fond of both and find them about equally desirable.

Children were tested with a comparable procedure adjusted to suit a classroom and human food preferences.

Results

Children in both groups were largely consistent, 18 of 20 choosing the container associated with an item *different* from the one the experimenter was eating.[1] That is, on trials on which the experimenter was seen eating a cookie, children selected the container which held the doughnut, and on trials on which he was seen eating the doughnut, selected the container which held the cookie.

Chimpanzees, however, were inconsistent. Sadie, the oldest chimpanzee, responded as did the children, choosing the container associated with a fruit *different* from the one eaten by the experimenter. She made this selection on the first trial as well as on all subsequent ones. By contrast, Luvie did the opposite, choosing the container associated with the fruit that was the *same* as the one eaten by the trainer. For instance, upon seeing the trainer eat the apple, she went to the container which held the apple, and upon seeing him eat the banana, to the container which held the banana. Bert and Jessie responded in an intermediate fashion, choosing the container associated with food *different* from that which the trainer was seen to be eating after first choosing the opposite container for two and four trials, respectively.

Discussion

Causal reasoning is difficult because a missing item must be reconstructed. While in simple learning, a monkey receives an electric shock when it presses a lever (and in observational learning observes that another receives an electric shock when it presses the lever), in causal reasoning the monkey does not observe the model press the lever but sees only its emotional response to the shock. Because in both simple and observational learning temporal contiguity between the lever press and electric shock are either experienced or observed, the relation between them is readily learned.

Causal reasoning, however, does not provide such temporally contiguous events. While the monkey sees the relation between the lever press and the model's painful state, the chimpanzee does not actually see this relation, but experiences only the informant's emotional state, and must reconstruct from it the event that caused the state. In our tests, even though the chimpanzees experienced the same emotional states as a consequence of the same events, they were incapable of reconstructing those events from the emotional state of another chimpanzee. This helps clarify the striking difference in difficulty between

[1]The difference between these data and preliminary data reported earlier (Premack, 1983) comes from an unaccountable difference between our village and city children. Village children typically lagged city children by 6–12 months.

learning and reasoning, and suggests why the former is found in all species, the latter in exceedingly few.

One might say that in the second experiment there is evidence for causal reasoning on the part of the children and perhaps one chimpanzee. This experiment can be seen as one of causal reasoning because here too there is a missing element to reconstruct. The subjects saw the trainer eating one or another food, but were never shown *where* he obtained it. Nevertheless, the children and perhaps one chimpanzee evidently did reconstruct the location of the food.

What is most interesting about this outcome is that subjects "asked" the question of where the trainer got the food, and "answered" it quite specifically by going consistently to the container holding the food different from that eaten by the trainer. They assumed the food was the same as that which had been placed in the container, and in making this assumption believed the one container to be empty.

It is not the specific content of the assumption alone that is of interest, but the additional fact that they made any assumption at all. Most species, we suggest, will make no assumptions. They will observe the trainer eating and never ask where he obtained the food. Such a question is asked only if one sees an event as part of a causal sequence in which there is a missing component.

Could we induce our subjects to change their assumption? Suppose there is insufficient time for the trainer to recover the food placed in the containers. Would this affect choice? Keeping all other conditions constant, we tested this possibility by wrapping the fruit and pastries in elaborate packages before placing them in the containers. Now the trainer could not possibly have obtained the food from the containers – there was not sufficient time for him to unwrap these items.

Children of 4 years and older were profoundly affected by this change. They no longer chose the container holding the pastry different from the one eaten by the trainer but chose at chance level between the two containers. By contrast, younger children and the chimpanzees were totally unaffected by the change in the availability of the item. They responded as before.

Labelling a causal action

A causal action can be analysed into three components: the actor who carries out the action, the object on which he acts, and the instrument with which he performs the action. For instance, John paints a wall with a brush, cuts an apple with a knife, and washes his dirty socks in soapy water. Can young children and chimpanzees analyse such causal sequences?

We devised a non-verbal procedure to answer this question by showing simple actions on a television monitor and giving our subjects markers that adhered to the screen and allowed them to identify each component of an action. A red square placed on John, for example, identified the actor; a green triangle on the apple, the recipient of the action; a blue circle on the knife, the instrument of the action.

Sarah was given this problem. She was trained on three different actions: Bill cutting an orange, John marking a paper with a pencil, and Henry washing an apple with water. The trainer demonstrated the proper placement of each marker, handed the markers to Sarah, and corrected Sarah's errors.

After reaching criterion on the three training tapes, she was given a transfer test in which all her responses were approved – our standard procedure in transfer tests. The tests were uniquely demanding, for the scenes were not merely new, but also decidedly more complex than those used in training.

Where the training scenes had presented one person acting on one object with one instrument, the transfer scenes presented: two objects, only one of which was acted upon; two instruments, only one of which was used; and two persons, only one of whom carried out the action, the other being engaged in some scenes as an observer of the action of the first, and in other scenes as the recipient of action; for example, Bill brushed Bob's hair.

Sarah passed the transfer tests, but at a relatively low level of accuracy. She was 85% correct in the use of the actor marker, 67% correct with the object marker, and 62% correct with the instrument marker. We attempted to improve her accuracy by training her on the transfer series, correcting her errors where previously we had approved all her responses. The attempt failed because she now placed the markers on the blank part of the screen (calling our attention to a fact we had previously overlooked – most of the screen is blank!). This tactic was entirely new, and brought the experiment to a halt.

In retrospect, we recognize the experiment was needlessly difficult, and could have been simplified by dropping one of the categories, either of object or of instrument.

Kim Dolgin (1981) as part of her doctoral research applied the same procedures to young children from 3.8 to 4.4 years, with an average age of 4 years, using the same non-verbal approach used with Sarah. The children failed the transfer tests. They did not properly identify the actor marker but drew a simpler distinction, animate/inanimate (or perhaps person/non-person). They reserved the actor marker for people but without regard for whether the person was an actor, observer, or recipient of the action. They made a similar simplification in the case of object and instrument markers, reserving them for non-persons, but without regard for whether the object or instrument was in actual use or simply present.

With the children Dolgin took a further step not possible with Sarah. She told the children the meaning of each marker. For example, she presented the scene in which Bill cut the apple, and then showing the child the actor marker told her

"He's the one doing it", the object marker "This is what he's doing it to", and the instrument marker "This is what he's doing it with." The results were dramatic. The children passed the transfer tests at an average level of 93%, far higher than Sarah.

Causal sequences as transformations

An actual causal sequence is a dynamic one in which an agent acts on an object, typically with the use of an instrument, changing its state and/or location. But one can represent the causal sequence in a stylized way – an apple and a cut apple representing the transformation of the apple, a knife the object responsible for the transformation.

To determine whether non-speaking creatures can recognize such transformations, we designed a test in which the subject was given an *incomplete* representation of this causal sequence and was required to choose an alternative that properly completed it (Premack, 1976).

The main actions we tested were cutting, marking, wetting and, in special cases, actions that reversed the effect: joining, erasing and drying. The subjects had extensive experience with the actions on which they were tested, carrying them out in a play context.

The problem was given to three chimpanzees (and numerous children) in two basic formats, one requiring that they choose the missing operator, another the missing terminal state. In the former, they were given as a representative test sequence "apple ? cut apple" along with the alternatives: knife, pencil, container of water. In the latter, they were given as a representative test sequence "apple knife ?" along with the alternatives: cut apple, cut banana, wet apple. The chimpanzees were given not only novel object–operator combinations but also anomalous ones – for example, cut sponges, wet paper, fruit that had been written on – and performed as well on the anomalous material as on the other (see Premack, 1976, pp. 4–7, 249–261 for details).

The tests were passed only by language-trained chimpanzees which were not given any training on the test but passed them on their first exposure to them. This test was only one of four that language-trained chimpanzees could do; the other three were analogies, same/different judgments on the relations between relations, and the matching of physically unlike proportions (e.g., $\frac{1}{4}$ potato to $\frac{1}{4}$ glass of water) (Premack, 1983). Language training conferred an immediate ability to solve this complete set of four tasks. Only with protracted training could non-language-trained chimpanzees be taught to do these tests, and then only one test at a time, with no apparent saving from one to the other (Premack, 1988).

Mapping the directionality of action

The direction of the action was depicted in the test sequences by the left/right order of the objects. Thus, the object in its initial state was always presented on the left, the object in its transformed state on the right. But the standard tests did not require the subjects to read the sequences from left to right; they might have chosen an operator simply on the grounds that it belonged with a particular action – a knife, for example, as the operator for cutting – making this choice without regard to order. Whether the intact apple was to the left or right of the cut one, the animal could well have made the same choice.

To obtain evidence that Sarah could discriminate the left–right order of the sequence and make use of it, she was first acquainted with pairs of actions that had reverse effects. For example, the trainer showed her how to mend broken pieces of an object with Scotch tape, how to erase pencil or crayon marks with a gum eraser, how to dry a wet object with a cloth – and then gave her the opportunity to carry out the actions. She adopted these new actions with enthusiasm.

She was then given three test sessions with the original "cut, mark, wet", and four sessions with the new cases "tape, erase, dry". The tests took the standard form, for example, "apple ? cut apple", accompanied by the standard three operators, for example, knife, container of water, pencil. A total of 26 objects were used and 12 operators, two of each kind. Each of the three cases (tape, erase, dry) was presented four times per session in random order, with the position of the correct operator randomized across left, centre and right positions.

Results:
Total = 40/60
Original cases = 12/18
New cases = 28/42
Z_{diff} between old and new not significant.

These preliminary results simply established that Sarah understood the new actions and could deal with them correctly. She was then required to use the left–right order of the sequence, and presented pairs of trials in which the same material appeared in reverse order. For example, "paper ? marked paper"; "marked paper ? paper". She was given pencil, container of water, and eraser as possible operators. Now, to choose correctly, Sarah had to take order into account, for while pencil was correct (and eraser incorrect) for one order, eraser was correct (and pencil incorrect) for the other. It takes a pencil to mark a blank sheet of paper, an eraser to remove the mark.

Sarah was given 16 sessions, 12 trials per session, old cases being presented 24 times each, new cases 36 times each in random order across sessions. Although the objects and operators used were the same as those in the previous tests, they were combined in new ways. All other details were the same as those of the preceding tests.

Results:
Total = 110/180
Original cases = 47/72
New cases = 63/108

If we exclude trials in which Sarah chose the incorrect irrelevant alternative rather than the relevant one, corresponding figures are:

Total = 110/148
Original cases = 47/62
New cases = 63/86
Z_{diff} between old and new not significant.

The reversal pairs compared as follows:

Cut/tape: 28/60
Mark/erase: 37/60
Wet/dry: 45/60
Z_{diff} between c/t and w/d = 3.18.

Finally, Sarah was given an extensive transfer test involving new objects and operators. In five sessions of 12 trials per session, 30 new objects were used as well as 60 new operators, 10 of each kind. Each object appeared twice, once with one or another of the six new operators, and again with the reverse operator. Each operator appeared three times: as correct alternative, incorrect reverse alternative, and incorrect irrelevant alternative. Each case appeared twice per session in counterbalanced order. All other details were the same as those already reported.

Results:
Total = 44/60
Original cases = 20/30
New cases = 24/30
Excluding trials in which Sarah chose incorrect alternative:
Total = 44/52
Original cases = 20/25

New cases = 24/27
Reversal pairs:
Cut/tape = 15/20
Mark/erase = 12/20 ($p < .05$ with three alternatives)
Wet/dry = 17/20
No significant Z_{diff}.

These data establish that Sarah could use left–right order to map the directionality of action as accurately on unfamiliar as on familiar cases.

Multiple transformations

The basic consequence of a causal action is a transformation – a change from an initial state to a final one. Could Sarah understand causal action from this perspective, looking at the initial state of an object, comparing it with its terminal state, then selecting the operator(s) that explains or accounts for the difference?

We can add to the interest of this question by removing the restrictions that were applied to the examples Sarah had been given. First, transformations involved more than a single action – for example, paper could be both cut and marked. Second, the initial state could be an already-transformed object rather than one in an intact or canonical state.

Now we not only lifted restrictions, but gave Sarah a special trash bin in which to discard incorrect or irrelevant operators. So, besides removing the interrogative particle and replacing it with the correct or relevant operators, she was required to select the incorrect or irrelevant operators, and place it/them in the trash.

The test consisted of six 12-trial sessions, each consisting of both single-action and double-action trials in equal number counterbalanced over the session. The six actions and their combinations were presented in equal number in each session counterbalanced over the session. The rest of the procedural details have already been reported.

Results:
Single transformations = 25/36
Double transformations = 24/36 ($p < .001$, both cases)
Total = 49/72

These results add to the evidence of Sarah's ability to use the test sequences as representations of action. Her analyses answered these implicit questions: (1)

What operator changed the object from its initial to its terminal state? (2) In applying this operator to this object, what terminal state did one produce? (3) Which operators caused the difference between the initial and terminal state, and which did not? In answering these questions, Sarah had to attend to the order of the test sequences, "reading" them from left to right.

We speculate that the representational ability required to pass these tests is that of a mind/brain which is capable of copying its own circuits. In carrying out an actual causal sequence, such as cutting an apple with a knife, an individual may form a neural circuit enabling him to carry out the act efficiently. But suppose he is not required to actually cut an apple, but is instead shown a representation of cutting – an incomplete depiction of the cutting sequence such as was given the chimpanzees – could he use the neural circuit to respond appropriately, that is, to complete the representation by choosing the missing element? Probably not, for the responses associated with the original circuit are those of actual cutting; they would not apply to repairing an incomplete representation of cutting. Moreover, the representation of a sequence can be distorted in a number of ways, not only by removing elements as in the chimpanzee test, but also by duplicating elements, misordering them, adding improper elements, or combinations of the above.

To restore distorted sequences to their canonical form requires an ability to respond flexibly, for example, to remove elements, add others, restore order and the like. Flexible novel responding of this kind is not likely to be associated with the original circuitry (that concerned with actual cutting), but more likely with a copy of the circuit. Copies of circuits are not tied, as are the original circuits, to a fixed set of responses, and they may therefore allow for greater novelty and flexibility. For this reason, we suggest, flexible responding may depend on the ability of a mind/brain to be able to make copies of its own circuits.

An attempt to combine three questions

Sarah was given a test that consisted of three questions: (1) What is the cause of this problem? (2) How can it be solved? (3) What is neither cause nor solution but merely an associate of the problem? These questions were not asked explicitly, with words, but implicitly with visual exemplars. Similarly, her answers were not given in words but in visual exemplars.

The problems about which she was queried were depicted by a brief videotape (the terminal image of which was put on hold). For instance, she was shown a trainer vigorously stamping out a small conflagration of burning paper. The questions asked her in this case were: What caused the fire? How could it be put out? What is neither cause nor solution but an associate of the fire?

The correct answers were photographs of: matches (cause), a bucket of water (solution), and a pencil (associate), the latter because she often used a pencil in scribbling on paper of exactly the kind that was shown in the videotape.

The three questions were identified by different markers (like those used to identify the three components of an action), though the meaning of these markers was not determined by their location on the television image, but by the correct answers with which each marker was associated.

The markers were introduced by presenting each of them with a videotape, offering three photographs and teaching her which of them was correct. In the example concerning a fire, which served as one of three training cases, she was presented the marker for "cause" (red square), with three photographs – matches, clay, knife – and taught to choose matches. When presented the marker for "solution" (green triangle), she was shown photographs of water, Scotch tape, eraser, and taught to choose the bucket of water. When presented the marker for associate (blue circle), she was shown photographs of pencil, apple, blanket, and taught to choose pencil. This procedure, repeated with two other training videotapes, was intended to teach her to view the problems depicted on the videotape according to the perspective indicated by the correct answer associated with a marker.

Once she reached criterion on the three training cases, she was given a transfer test involving 20 new problems. When she failed, it was decided to train her on this material and to bring in a new trainer – one who no longer played an active role in her care or testing but who had been an early caretaker, was a favourite, and could be counted on to bring out her best effort. Sarah definitely "tried" harder with some trainers than with others. Ordinarily she looked only once at an alternative before choosing, but with a difficult question and a favourite trainer, several times.

Her looking behaviour was readily observable; after the trainer gave her the test material in a manila envelope, he left the cage area. Sarah could be observed on a television monitor to empty the envelope on the cage floor, spread out the alternatives, inspect them, choose, and then ring her bell (as a period marker signalling an end and summoning the trainer).

With this favourite trainer she not only looked at the alternatives with more than usual care but did several double-takes; that is, looked, looked away, and then quickly looked back. This did not help her cause; she made three consecutive errors. When the trainer entered to show her the fourth videotape, she lost sphincter control and ran screaming about the cage.

Although a demanding test, it is not necessarily beyond the capacity of the chimpanzee. It must be taught more carefully than we did, not as a combination of three markers, but one marker at a time, then two and, only when there is success on two, all three presented together. We subsequently used this approach with $4\frac{1}{2}$–$6\frac{1}{2}$-year-old children on a problem only slightly less demanding than the one given Sarah, and they succeeded nicely (Premack, 1988).

General discussion

There are two traditions in which causality has been studied in psychology: the natural causality of Michotte (1963), and the arbitrary causality of Hume (1952).

Natural causality (Premack, 1991) concerns the relation between special pairs of events – one object launching another by colliding with it is the classic example; whereas *arbitrary causality* concerns the relation between any pair of temporally contiguous events – a lever press followed by the presentation of a Mars bar is one example. These two traditions have fostered conflicting interpretations of causality.

The perception of natural causal relations requires only a single pairing of appropriate events, is demonstrable in 6-month-old infants (e.g., Leslie & Keeble, 1987) and is considered innate; whereas the perception of arbitrary causal relations requires repeated pairings in time of two events, and is learned. But are these differences real or do they simply reflect differences in subject matter?

Although Michotte's case is typically the only cited example of natural causality, it is essential to recognize that there is another, and more basic example of natural causality. This is the psychological case where we perceive causality under two conditions: (1) when an object moves *without* being acted upon by another object, that is, is "self-propelled" (Premack, 1990); and (2) when one object affects another "at a distance", that is, affects another despite a lack of spatial/temporal contiguity.

Humans unquestionably perceive causality under both these conditions. Yet this fact has received little comment – virtually none compared to the extensive comment on the Michotte case. Why? We suggest, because the perception of causality in the psychological domain evolved far earlier than it did in the Michotte case and belongs to a "part of the mind" that is less accessible to language.

The perception of causality of the Michotte variety probably evolved late, and even then only in a few tool-using species (Kummer, in press), for only humans, apes, and the exceptional monkey (e.g., *Cebus*) handle objects in a manner capable of producing collisions. Compare this to the perception of causality in the psychological domain, which is not restricted to a few tool-using species but is found in virtually all species. Intentional action which involves either a single individual or one individual acting upon another is part of the experience of all but a few invertebrates.

Bar pressing that produces food, threats that produce withdrawal, collisions that launch objects, and so on, all fit neatly into either a physical or psychological category. These cases are important because they give the impression that the concept of cause has a content: psychological, physical, or both. However, infants may perceive a causal relationship when presented with totally arbitrary cases; for example, a sharp sound followed by a temporally contiguous colour change in an object. If so, this example is important because it demonstrates that the concept of cause may be without content.

Let us habituate one group of infants to a contiguous case, another to a delay case, and then apply the Leslie–Keeble paradigm by reversing the order of the two events. We present colour change *followed* by sound to both groups. The contiguous group, we predict, will show greater dishabituation (recovery in looking time) than the delay group.

In other words, our outcome will be the same as that obtained by Leslie–Keeble in the Michotte example – a greater recovery for the group in which the events are presented contiguously. But how does one explain these results? Just as Leslie–Keeble do. While contiguous events certainly do lead to the perception of causality, is this perception confined to the Michotte case?

Causality is not bound by content, we suggest, for the sequence "sound–colour change" which involves neither an intentional act nor the transfer of energy from one object to another, is an example of neither psychological nor physical causality. Perhaps the concept of causality at its most fundamental level is no more than a device that records the occurrence of contiguous events – associative learning – and is found in all species.

All species may share a device that records the occurrence of contiguous events, and evolution contributed two major additions to this primitive device: first, the capacity to act intentionally which enabled certain species not only to register but also to produce contiguous events; second, the capacity, largely unique to the human, to explain or interpret events that have been both registered and produced.

While the basis of the primitive level has not been resolved by neuroscience (for this level operates on "events", and how the mind/brain *binds* items so as to construct events remains a challenge for neuroscience; e.g., Singer, 1990), fortunately, the second level of causality is well represented by work on animal learning. Especially Dickinson and his colleagues (e.g., Dickinson & Shanks, in press) have considered the special subset of concurrent events – act–outcome pairs – brought about by intentional action. Are such pairs marked in some fashion, and thus represented differently in memory from other concurrent pairs?

The third level of causality is to be found in recent work on domain-specific theories of naive physics (Spelke, Phillips, & Woodward, in press; Baillargeon, Kotovsky, & Needham, in press), psychology (Leslie, in press; Gelman, Durgin, & Kaufman, in press; Premack & Premack, in press), and arguably biology (Carey, in press; Keil, in press). These theories separate the concurrent events (registered by the primitive device) into special categories, and propose an explanatory mechanism for each of them. Explanation, embedded in naive theories about the world, is largely a human specialization.

References

Baillargeon, R., Kotovsky, L., & Needham, A. (in press). The acquisition of physical knowledge in infancy. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Carey, S. (in press). On the origin of causal understanding. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Dickinson, A. & Shanks, D. (in press). Instrumental action and causal representation. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Dolgin, K.G. (1981). *A developmental study of cognitive predisposition: A study of the relative salience of form and function in adult and four-year-old subjects*. Dissertation, University of Pennsylvania.

Gelman, R., Durgin, F., & Kaufman, L. (in press). Distinguishing between animates and inanimates: Not by motion alone. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Gillan, D.J. (1981). Reasoning in the chimpanzee. II. Transitive inference. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 150–164.

Gillan, D.J., Premack, D., & Woodruff, G. (1981). Reasoning in the chimpanzee. I. Analogical reasoning. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 1–17.

Hume, D. (1952). An enquiry concerning human understanding. In *Great books of the western world* (Vol. 35). Chicago: Benton.

Keil, F. (in press). The growth of causal understandings of natural kinds. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Kummer, H. (in press). On causal knowledge in animals. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Leslie, A., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, *25*, 265–288.

Michotte, A. (1963). *The perception of causality*. London: Methuen.

Premack, D. (1976). *Intelligence in ape and man*. Hillsdale, NJ: Erlbaum.

Premack, D. (1983). The codes of man and beasts. *Behavioral and Brain Sciences*, 6, 125–167.

Premack, D. (1988). Minds with and without language. In L. Weiskrantz (Ed.), *Thought without language*. Oxford: Clarendon Press.

Premack, D. (1990). The infant's theory of self-propelled objects. *Cognition*, *36*, 1–16.

Premack, D. (1991). *Cause/induced motion: intention/spontaneous motion*. Talk at Fyssen Foundation Conference on the origins of the human brain. Clarendon Press: Oxford.

Premack, D., & Premack, A.J. (in press). Semantics of action. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.

Singer, W. (1990). Search for coherence: A basic principle of cortical self-organization. *Concepts in Neuroscience*, *1*, 1–26.

Spelke, E.S., Phillips, A., & Woodward, A.L. (in press). Infant's knowledge of object motion and human action. In A.J. Premack, D. Premack, & D. Sperber (Eds.), *Causal cognition: A multidisciplinary debate*. Oxford: Clarendon Press.