

DOCUMENT RESUME

ED 297 882

PS 017 522

AUTHOR Chi, Michelene T. H.; And Others
TITLE Knowledge-Constrained Inferences about New Domain-Related Concepts: Contrasting Expert and Novice Children.
INSTITUTION Pittsburgh Univ., Pa. Learning Research and Development Center.
PUB DATE Mar 88
NOTE 40p.
PUB TYPE Reports - Research/Technical (143)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Abstract Reasoning; *Classification; *Cognitive Ability; *Cognitive Structures; Comparative Analysis; *Knowledge Level; Primary Education; *Young Children
IDENTIFIERS *Concept Acquisition; Dinosaurs; Rationality; Representational Thinking

ABSTRACT

Three studies examined the domain of concepts about dinosaurs in order to assess how the domain might be structured in 4- through 7-year-old children's representations and to explore how the knowledge might be used. Findings indicated that significant differences exist in the way expert and novice children's representations are structured. Evidence further suggested that the structure of expert children's knowledge is more coherent, both locally and hierarchically. Some experts were able to sort dinosaurs at both the superordinate and family levels, but no novices did. Expert children could discriminate contrastive diet classes by using a single feature in both an inclusive and exclusive way. Expert children could reason categorically, using superordinate, family, and dinosaur categories. Novice children were just as competent as matched expert children in using general learning skills when these skills were assessed in a domain in which both the expert and novice children had equivalent knowledge. The use of learning skills seemed to differentiate experts from novices only in the context of dinosaur knowledge. Findings suggest that children can act more or less intelligently on the basis of what knowledge they have. A modest mastery of a domain results in intelligent explanations, constrained inferences, categorical reasoning, hierarchical classification, and classification based on well-defined family structures. Six figures are presented. (RH)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

ED297882

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

This document has been reproduced as
received from the person or organization
originating it.
 Minor changes have been made to improve
reproduction quality.

• Points of view or opinions stated in this docu-
ment do not necessarily represent official
OERI position or policy.

Knowledge-Constrained Inferences About New Domain-Related Concepts: Contrasting Expert and Novice Children

Micheline T. H. Chi, Jean E. Hutchinson,
and Anne F. Robin



Learning Research and Development Center
University of Pittsburgh
Pittsburgh, PA 15260

1988

"PERMISSION TO REPRODUCE THIS
MATERIAL HAS BEEN GRANTED BY

Pittsburgh Univ.
(LRDC)

TO THE EDUCATIONAL RESOURCES
INFORMATION CENTER (ERIC)."

PS 017522

**Knowledge-Constrained Inferences
About New Domain-Related Concepts:
Contrasting Expert and Novice Children**

Michelene T. H. Chi

Jean E. Hutchinson

Anne F. Robin

University of Pittsburgh

Running Head: KNOWLEDGE-CONSTRAINED INFERENCE

March 1988

Abstract

This paper discusses research from three different studies on domain-specific knowledge, with dinosaurs as the specific domain. The three studies focus on two related issues: first, to define "structure" as it applies to a conceptual domain (i.e., what it means to say that the organization of dinosaur knowledge is more or less structured); and second, to relate the way the knowledge is structured to how it is used. In particular, we explore ways in which the structure of existing knowledge constrains our learning of new domain-related concepts.

**Knowledge-Constrained Inferences
About New Domain-Related Concepts:
Contrasting Expert and Novice Children**

A central concern for developmental psychologists is the issue of structural change. Do knowledge structures undergo major changes as children mature? According to Piaget, structural change is "global" in the sense that children at each developmental stage are supposed to have acquired a particular representational structure within which they can reason logically across all domains. The particular structure that children are capable of reasoning within is represented by the logical inferences that children can engage in, such as transitive inferences, reversible operations, and so forth. In a sense, the structures in Piaget's representations are operators that can act on any domain of knowledge. Thus, when children have reached the stage in which they are capable of transitive inferences, they should be able to infer in a transitive way for all domains. Since these structures constrain children's ability to make logical inferences across all domains, restructuring in cognitive development constitutes a global change.

It is becoming increasingly clear that Piaget's notion of global restructuring is too general. In many instances, having knowledge in a specific domain can overcome any limitations that could have been imposed by the lack of global operators. On the other hand, lacking knowledge in a specific domain can also prevent adults from reasoning logically, even though they are presumed to have the logical operators. (For summaries of this kind of evidence, see review papers by Gelman & Baillargeon, 1983, on Piagetian tasks, and Chi & Ceci, 1987, on memory-related tasks.)

The dilemma created by the role of domain knowledge necessitates the examination of the following issues:

1. what is the structure of a domain of knowledge;
2. what does it mean to say that knowledge of a domain becomes more structured;
3. how does the structure of domain knowledge relate to the child's ability to remember and reason logically;
4. and what are the mechanisms that produce the greater structure.

This paper focuses on all of these issues to some degree, but more so on the first three issues.

How Domain Knowledge Is Represented

In order to answer questions 1 and 2 above about how domain knowledge is structured, one must first create a representation of the domain knowledge. There are many ways to represent knowledge of static concepts such as dinosaurs. A popular way is in terms of a network of nodes and links. A node usually corresponds to any concept that we designate, and a link corresponds to the relation between nodes. However, very little *empirical* work has been done to define what it means to say that one network is more or less structured than another. Our first study

(Chi & Koeske, 1983) focused primarily on how to describe the degree of structure that exists within a domain of knowledge.

When we discuss the degree of structure within a representation of static concepts, we are not referring primarily to the quantity of knowledge. Although the amount of knowledge necessarily determines to some degree the structure of the knowledge, it is important to unconfound the two factors, if possible. We are most concerned with those characteristics of a structure that can effect the use of the knowledge. Hence, the kinds of representational properties that we predict are important are the *coherence* of a structure and the *hierarchical* nature of a particular structure. Coherence refers to both the integratedness of the whole structure (henceforth called *global coherence*) and the integratedness of the substructures that form the hierarchy (*local coherence*). In this paper, because we are talking about concepts that have a clear classification scheme, global coherence will refer to the presence of hierarchical organization. We will describe what constitutes evidence for coherence and hierarchy below.

We postulate that two sets of identical concepts and attributes can have different degrees of local and global coherence simply by the way the concepts are linked. In Figure 1, for example, both the top and the bottom networks have identical concepts (the circles) and attributes (the triangles), and the same number of links. However, because the concepts and attributes are linked in different ways, we claim that the behavioral manifestations corresponding to the two networks will be different. Thus, one could say that Structure I is "better" or "more structured" than Structure II. In a sense, in these hypothetical networks, we have controlled for the content and amount of knowledge.

Insert Figure 1 about here

Within these hypothetical networks, we can thus define *local coherence* of the substructures in terms of the patterns of interlinking as well as attribute-sharing among concepts, and *hierarchy* (or *global coherence*) in terms of the patterns of relationships among substructures. Comparing Structures I and II of Fig. 1, we might say that Structure I is more globally as well as more locally coherent than Structure II. Structure I is more hierarchical than Structure II because it has well-differentiated substructures which together constitute the larger structure (Concepts A and C form one substructure; Concepts B and D form another). Structure II does not have any hierarchical embedding. The substructures in Structure I are also more locally coherent in the sense that the pairs of concepts

share attributes to a greater degree than any subsets of concepts in Structure II. Thus, one might say that Concepts A and C in Structure I are more related than Concepts A and C in Structure II. (Such relatedness can be manifested in clustered recall, for example.) Of course, Structure II might be considered to be more locally coherent in the context of a larger structure. Thus, coherence and hierarchy are necessarily relative terms.

One way to describe the relationship between local and global coherence is in terms of *differentiation*. Structure I is more differentiated than Structure II in that it has substructures that share more attributes within themselves than they do between them. One can increase the differentiation of Structure II and thereby cause it to become more hierarchical if one adds direct and indirect links between Concepts A and C and between Concepts B and D to form coherent substructures (see Structure IIA, Figure 2). One can also decrease the differentiation of Structure I and cause it to become less hierarchical if one adds links from the concepts and attributes of one substructure to the concepts and attributes of the other substructure (see Structure IA, Figure 2).

Insert Figure 2 about here

Our notion that coherence and hierarchy may be a fundamental way of defining structure are borne out by our data. In our first study (Chi & Koeske, 1983), we elicited a 4 1/2-year-old child's knowledge of dinosaur concepts by asking him: (a) to generate a sequence of dinosaur names, (b) to identify a dinosaur from dinosaur attributes that we provided and, (c) to generate dinosaur attributes when given dinosaur names. The sequential generation of dinosaur names provided dinosaur-dinosaur linkages. The identification of a dinosaur from its attributes, as well as the generation of attributes from dinosaur names, provided the data for the dinosaur-attribute links. The only assumption that was made in mapping the network was that no redundant nodes were depicted. Thus, if the child generated the attribute "sharp teeth" for two different dinosaurs, then these two different dinosaurs were shown to be sharing the same "sharp teeth" node.

For simplification and comparison purposes, we divided the total set of dinosaurs that the child generated into 2 subsets of 20, on the basis of how familiar the child was with each of the dinosaurs. The classification into each of the two categories (Better-Known and Lesser-Known) was done on the basis of two external sources: the mother's judgment about which dinosaurs she thought her child knew more or less about, and the frequency with which a given dinosaur was mentioned in the books that were read to the child. Each group of 20 dinosaurs could also be divided into 7 subgroups on the basis of their family membership, as discussed in the dinosaur books.

For further simplification, the networks that we drew for the better- and lesser-known dinosaurs consisted of the attributes and links of 7 dinosaurs from each subset of 20. These seven dinosaurs (the targets) were chosen so that they could be matched on the amount of information (number of attributes) generated by the child about each dinosaur. We wanted to control for the number of attributes generated because we were not interested in the effect produced by having a greater quantity of knowledge. In this case, the child knew about five attributes for each of the seven target dinosaurs from both the better-known and lesser-known sets. However, although the two subsets of seven targets were matched on the number of attributes known about each, each subset of seven belonged to a different number of families. The seven better-known target dinosaurs belonged to only two families--armored dinosaurs and giant plant eaters--whereas the seven lesser-known target dinosaurs belonged to five families: armored dinosaurs, small bird eaters, water dwellers, duckbills, and early meateaters.

Once we had represented the target dinosaurs in terms of their links and attributes in network form, we could quantify coherence and hierarchy in an operationally definable way. Basically, our networks corresponded to the hypothetical ones depicted in Figure 1. The network of the seven better-known dinosaurs (see Fig. 3) corresponded to the network depicted in Structure I, with two families as substructures. Target dinosaurs were indicated by rectangles and families were indicated by dotted circles in Fig. 3. Families in our network were defined by clusters of interrelated concepts and attributes. A family is not represented by a single node, as is often done in simple hierarchical representation of concepts. The network of the lesser-known dinosaurs (shown in Fig. 4) corresponded to the network depicted in Structure II, but with potentially five substructures corresponding to the five families.

Insert Figures 3 and 4 about here

For the better-known dinosaur network in Fig. 3, then, evidence of global coherence can be gleaned from the existence of multiple links among target dinosaurs within the same family, but only single links to target dinosaurs from other families. Evidence of local coherence can be seen in the great degree of sharing of attributes among dinosaurs of the same family. This pattern of greater within- and lesser between-family linkages was evident for the entire set of 20 better-known dinosaurs as well, not just the 7 target ones. The pattern of linkages for the lesser-known dinosaurs was more diffuse and less differentiated. There was no sharp contrast between the linkages within and the linkages between families. Like Structure II of Fig. 1, the connections for the lesser-known dinosaurs were much more uniformly distributed among the families. Thus, we concluded that there was no evidence of local or global coherence in the network of the lesser-known dinosaurs.

In this study, we also had three additional measures of the child's dinosaur knowledge: the child was asked to (a) categorize the 40 dinosaurs in a free sorting task, (b) free recall each list of 20 dinosaurs after it was read out-loud to him and (c) a year later, name each of the 40 dinosaurs when a picture of it was shown. In the free recall task, the child could recall around 45% of the dinosaurs taken from the better-known set, and only 22% of the dinosaurs taken from the lesser-known set. Furthermore, the recall from the better-known list had a higher clustering score ($RR=.67$, where RR is Bousfield's (1953) ratio of repetition) than the clustering score of the lesser-known set ($RR=.17$). We interpret the better recall and clustering for the better-known set as validating our analyses of the differential coherence and hierarchy of representations as a function of knowledge. The 20 better-known dinosaurs must have had greater local coherence (more shared links within each family) and more differentiated family structures in order to produce a higher clustering score and better recall. The 20 dinosaurs of the lesser-known set may not have had as cohesive a pattern of linkages. The fact that after a year, the child could name many more dinosaurs from the better-known than the lesser known set (11/20 vs. 2/20) also suggests that the attributes required for identification are not as strongly associated for the lesser-known dinosaurs. (This distinction did not show up a year earlier, when the child could identify all 40 dinosaurs by name.)

The results from the sorting task were the most intriguing because they were not predictable from the network representation. (The sorting data were not included in the Chi & Koeske (1983) paper because we worried at the time that it was not a robust result. We now have data from two other studies, reported later in this paper, that replicate this finding.) The child sorted the 20 better-known dinosaurs very quickly (without any hesitation or pauses) into two groups: meat-eaters and plant-eaters. This pattern of sorting was consistent across two separate trials. In contrast, he sorted the 20 lesser-known dinosaurs with a great deal of hesitancy and uncertainty. Furthermore, in three trials, he could not reach a consistent sort. The first sort was divided between plant- and meat-eaters, the second sort combined meat-eaters with land-dwellers, and included a new group of water-dwellers, and the third sort included land-planteaters, water-planteaters, and water-meateaters. Thus, the child's sorting behavior for the better-known dinosaurs resembled that of older children's sorting behavior in general, in that it was very fast, exhaustive, and stable across trials. His sorting for the lesser-known dinosaurs resembled the behavior of younger children: it was hesitant, inconsistent, and slow.

The dilemma we had about his sorting behavior was that the way the child sorted the dinosaurs did not map clearly onto the family structures that we had represented in the networks, as well as the family structures that were manifested in the clustering pattern of the free recall. This suggests that there is a higher level of the hierarchy that we did not capture: that is, the families probably coalesced into two superordinate groups--meat- and plant-eaters. The more consistent sorting of the better-known dinosaurs could be interpreted to mean that the families of these 20

dinosaurs have coalesced clearly into these two groups. Had we represented all 20 of the better-known dinosaurs in a network, we might then have seen how the families related to each other in a hierarchical way.

We can speculate on how children coalesce families into higher-level groups that contrast with each other. Children may come to realize that there exist one or more features that most of the members of one set of families possess, and most of the members of a contrasting set of families do not possess. If we re-examine the network for the better-known dinosaurs (see Fig. 3), we see that most of the target dinosaurs have links to the same diet attribute, in this case planteating (as represented by "di," and highlighted here by shading). We now can re-interpret these linkages to mean that this is the critical attribute which one family (the armored dinosaurs) shares with the other family (the giant planteaters). That is, we can think of the "planteating diet" node as the main attribute that is shared by the two families of dinosaurs depicted in Fig. 3. This is perhaps the way that hierarchical groups are formed.

In contrast to the network for better-known dinosaurs, the network for the lesser-known dinosaurs shown in Fig. 4 depicts many "di" nodes. These nodes correspond to the mention of different diets associated with different dinosaurs. For example, some dinosaurs were said to eat eggs, others to eat meat, and still others to eat mushy plants. The child clearly had not generalized the similarities of the diet across dinosaurs, and did not see the families as sharing the same diet. Furthermore, many of the lesser-known dinosaurs' diets were incorrect and inconsistent. For example, ornithomimus is indicated to eat plants as well as eat eggs, and plateosaurus is indicated to eat meat as well as eat plants. These inconsistencies must also make it difficult to clearly segregate the families in the child's representation of the lesser-known dinosaurs. The presence of both lack of generalization and inconsistency suggests that well-defined hierarchical clustering of families has not emerged for the lesser-known dinosaurs.

Relationship Between Structure and Use of Knowledge

The major focus of the first study was to define a way of representing declarative knowledge of concepts (such as dinosaurs). In order to give some validity to the representations that we have depicted, we took several measures (recall, retention, and sorting) on a single child. The results of these measures all corresponded to some degree to the distinction we make between more- or less-structured knowledge. The distinction was based primarily on the existence of links among concepts, rather than the existence of concept nodes.

Our second study (Gobbo & Chi, 1986) added greater support for the idea that a more locally coherent and integrated knowledge structure is one in which there are more connections or stronger links among the concepts and attributes within a family. In this study we also explored the effects of knowledge structure on the use of that knowledge. The study compared five 7-year-old children who were experts in the domain of dinosaurs with five

7-year-olds who were novices. Expertise was defined on the basis of a pretest of dinosaur knowledge.

In one task, children were asked to sort the 20 pictures of better-known dinosaurs used in Study 1. These results will be reported in conjunction with the results from the third study, since the same task was also presented to additional subjects in that study. We had two other measures to evaluate the cohesiveness of a representation. First, we found that when children were simply asked to generate comments about dinosaurs, the expert children's protocols contained more connecting words (such as "because", "if") than the protocols of the novice children. The greater presence of these connecting words in the expert children's protocols suggests that when they activate a dinosaur concept node, several other related attributes and concepts also get activated with high strength, and thus, they feel compelled to continue and state them. The connecting words serve the function of accessing additional knowledge that is associated with the information that is presented in the dinosaur picture. In the following quote:

"And he had webbed feet, so that he could swim, and his nose was shaped like a duck's bill, which gave him his name".

the propositions generated prior to the connecting words (webbed feet, duckbill-like nose) were derived from the explicit external features depicted in the picture, whereas the propositions generated after the connecting words (he could swim, gave him his name), were derived from implicit knowledge associated with those features.

In contrast, novice children's productions tended to be a listing of explicit features that were depicted in the picture of the dinosaur, such as:

"He has sharp teeth. He has three fingers.

He has sharp fingers, sharp toes, a big tail."

Because novice children lack a strongly interconnected set of attributes and concepts, their protocols do not contain a great number of connecting words, nor do they contain inferences derived from the explicit external features. Consistent with this interpretation is the second finding, that expert children are much more likely than novices to change topics in their discussion of each dinosaur. This suggests that each topic (such as the manner of defense of a dinosaur) is strongly related to other features of the dinosaur (such as its diet, its habitat, its social behavior, and its physical appearance) in the child's representation. For example, the way a dinosaur defends itself depends on where it lives. If it lives in the swamp, it might hide in the water for defense. How the dinosaur defends itself also relates to the kind of protective armor that it has (its appearance). It is easier to understand the attributes of a dinosaur if one knows how they relate in a causal or correlated structure. We suggest that expert children have represented these

causal and correlated associations, so that it is difficult for them to discuss one theme without relating it to another. This results in the greater frequency of transition among the different themes.

What is the effect of the structure of children's knowledge on their use of that knowledge? In this second study, we examined usage of knowledge in a less traditional way (as opposed to the more traditional measures of recall and retention used in the first study). One effect of a more integrated and highly associated knowledge structure is hinted at in the analysis of the use of connecting words. That is, expert children, when confronted with a picture of a familiar dinosaur, often mention attributes of that dinosaur that are not pictorially depicted (implicit attributes). This is because a visual depiction of the dinosaur triggers other associated attributes and knowledge of that dinosaur. This is seen even more clearly in a propositional analysis of the children's protocols. Both the expert and novice children stated the same total number of propositions concerning the explicit features of a dinosaur (such as the sharpness of its teeth), but the expert children stated four times as many propositions concerning implicit features of a dinosaur (such as where it might live, or how it might attack an enemy).

Another major difference in the way expert and novice children use their knowledge lies in the number of comparisons they make among dinosaurs. Experts use comparisons more frequently than novices. These comparisons include instances in which children contrast dinosaurs as well as mention similarities. For example, when two dinosaurs share common features, expert children might generate comparative comments such as:

"They are pretty much like Diplodocus"

"They eat plants, kind of like Stegosaurus".

On the other hand, they are also capable of contrasting dinosaurs:

"And it goes in the water, but it doesn't go in like Brontosaurus".

These two uses of comparisons by experts suggests that their dinosaur concepts are both more interrelated, as well as more differentiated, than those of novices.

How knowledge is used and represented is further revealed by the way children decide upon the diet of a dinosaur. Since diet is a feature that most children mention, we wanted to see what features children use to derive the diet information. We found that expert children use predominantly four features. Furthermore, they can discriminate contrastive diet classes by using a single feature in both an inclusive and exclusive way. That is, if a feature is present, then the diet is meat-eating, but if the feature is absent, then the diet is plant-eating (or vice versa).

The presence and absence of a single feature correlated with diet may have been the factor leading experts to subcategorize dinosaurs into complementary diet classes. We speculate that the formation of hierarchical clusters is based on the development of these contrastive discriminations. This is consistent with our speculation from Study 1 that families coalesce into higher order categories if members of all the families share similar attributes. (See Chi, in press, for other suggestive evidence.)

Novices did not use features in such a contrastive way. They had different sets of features to determine whether a dinosaur was meat-eating or plant-eating. Consequently, the total number of features used by the novice children doubled those used by the expert children (9 vs. 4). This suggests that for novices, the two diet classes are not complementary, and thus are not well-differentiated. This is also evident in the expert child's representation of the lesser-known dinosaurs in Study 1 (Fig. 4). There, he would sometimes mention contradictory diets for a given dinosaur. Thus, the representation of novices is similar to the expert child's representation for the lesser known familiar dinosaurs.

In summary, the results of Study 2 provide additional evidence to suggest that the experts' representation of families is locally coherent. This is shown, for example, by the frequent use of comparisons among dinosaurs, of the same family. The results of this study also indicate the greater integratedness of knowledge about individual dinosaurs for the experts. This is shown by the way connecting words are used to draw out knowledge of a dinosaur's implicit attributes, as well as by the way experts frequently change topics in their discussion of a particular dinosaur. At the same time, experts' representations are also more differentiated than the novices, as evidenced by their use of a single set of features to contrast two diet classes, suggesting that they can combine families to form the higher-level categories of plant- and meateaters. Thus, the experts' representation of dinosaur knowledge seems to be more hierarchical and locally coherent than the novices' representation.

How Knowledge is Used to Learn New Concepts

In the third and just completed study of this series, we pursued the understanding of the same two issues (namely, to define "structure" as it applies to a conceptual domain, and to relate the way the knowledge is structured to how it is used), but with three additional specific goals. First, we wanted to see how knowledge is used to constrain learning about new domain-related concepts; second, we wanted a clearer conception of the nature of the knowledge that is embodied in a family (or more broadly speaking, a schema); and finally, we wanted to control more carefully our selection of the subject sample, so that they were matched on all dimensions except for their knowledge about dinosaurs. Because this third study has not been previously published, we report it in greater detail than the first two studies.

Two groups of 5 children were selected as subjects on the basis of their dinosaur knowledge, which was assessed by a pretest (the same as that used in Study 2). The expert group scored a mean of 59%, and the novice group scored a mean of 13% on the pretest. Each group of children included 3 boys and 2 girls, and the mean ages were 5 years 11 months for the expert group, and 6 years 3 months for the novice group. The two groups of children were further matched on a number of other measures, such as their forward and backward digit spans, and their scores on the Raven's Colored Progressive Matrices and the Peabody Picture Vocabulary test.

The stimulus materials consisted of line drawings of a set of 20 *familiar* dinosaur illustrations (the same as the better-known dinosaurs used in Studies 1 and 2), and a set of 20 *novel* dinosaur illustrations. The *familiar* set was composed of prototypical members of five taxonomic dinosaur families: duckbills, giant planteaters, horned dinosaurs, armored dinosaurs, and early meateaters. Members of a dinosaur family share a common set of explicit attributes (i.e., physical features) and implicit attributes (i.e., diet, habitat). These familiar dinosaurs were ones that both experts and novices had seen in common children's dinosaur books.

The *novel* set included only dinosaurs which were novel to both the experts and novices (i.e., were not commonly found in children's dinosaur books). Five of the stimuli were *prototypical* (yet novel) members of the five dinosaur families described above. Ten others were constructed, by physically transposing parts of dinosaur illustrations, to represent five *peripheral* family members (one from each family) and five *conflicting* family members. A *peripheral* dinosaur was composed mostly of features specific to one family with the exception of one conflicting feature that was salient to another family. Members of the *conflicting* set, on the other hand, were made up of physical features characteristic of several families. The remaining five stimuli, named *newly-discovered*, were recently discovered dinosaurs that did not belong to any of the five common families. These dinosaurs can be found in recently published dinosaur books.

The experimenter met with each child for approximately 4 one-half hour sessions, over a 2- to 3- week period, to complete the following five tasks (in the following fixed order). In the Sort With Familiar Dinosaurs Task, the 20 *familiar* dinosaurs were randomly displayed on a table in front of the child. Instructions were to "Put the dinosaurs that you think belong together in the same group." If a child constructed very large groups, she was asked to break them down further. Upon completion of the sorting, the child was asked to explain each grouping. No limit was placed on either the number or size of the groups to be formed. This task should indicate to what extent experts and novices categorize dinosaurs on the basis of family distinctions, as well as the differential use of more perceptually based classification.

In the Free Generation Task, 8 of the *novel* stimuli, randomly selected for each child so that two were chosen from each category (i.e., *prototypical*, *peripheral*, *conflicting*, & *newly-discovered*), were presented in random order with the instructions "Tell me the name of this dinosaur and everything you know about it." The purpose of this task was to assess how children infer information about novel domain-related concepts.

In the Probe Question Task, children were probed with specific questions, in case they were reticent and did not spontaneously generate many remarks in the Free Generation task. For each of the remaining 12 *novel* stimuli not used in the Free Generation task, children were asked the following five probe questions:

1. What would this dinosaur eat? Why?
2. Describe the place this dinosaur might live? Why?
3. What would this dinosaur's cousin be like? Why?
4. How did this dinosaur defend himself? Why?
5. If this dinosaur only ate plants/meat, would it be a peaceful dinosaur or would it attack other dinosaurs? Why?

The purpose of this task was to probe the degree to which children have integrated and related knowledge about a family of dinosaurs. For example, if they know about the diet of a dinosaur, they may have classified the dinosaur into either a meat-eating or a plant-eating class. If so, then any categorical knowledge about the meat- or plant-eating classes should constrain additional inferences that they might make about other attributes of the dinosaur. This task attempts to assess this kind of integrated knowledge.

In the Forced Choice Sort, the original groups into which the *familiar* dinosaurs were initially sorted by each child were spread out on a table. Then the child was asked to sort the set of *novel* stimuli either into these existing groups or into new groups. The purpose of this task was to examine the extent to which children perceived the *novel* dinosaurs as fitting into the family groupings of the *familiar* dinosaurs.

Finally, in an Oddity Task, the experimenter displayed 10 sets of four *familiar* dinosaurs one set at a time, saying "Three of these dinosaurs are alike in some way. The other one is different. Which one is different from the others? Why?" Within each set, three dinosaurs could be considered to be alike on a perceptual basis (i.e., stand on two legs) and three dinosaurs were similar on a conceptual basis (i.e., eat plants). This task should indicate which dimensions (i.e., physical attributes, diet, habitat) the child uses to categorize dinosaurs. Our hypothesis was that the experts would be more likely to categorize on the basis of conceptual features, whereas novices would be more likely to categorize on the basis of explicit physical features.

Results from the Free Generation Task

The free generation task was in a sense the purest measure of how children use knowledge to learn new concepts, because it was the first task utilizing the novel stimuli, no information was provided by the experimenter, and children were free to say anything they liked about the novel dinosaurs. We analyzed the free generation task in several different ways. These analyses will serve to illustrate some general themes that will then be discussed in relation to the other tasks.

Kinds of knowledge used to generate inferences. We began by breaking the protocols produced by the children down into propositions, or the smallest possible meaningful statements (see examples in text below). Experts and novices produced about the same total number of propositions across the eight novel dinosaurs (about 66 propositions). This suggests that we can be confident that both the novices and the experts try equally hard to generate inferences about novel dinosaurs. Most of the propositions could be classified as either Implicit or Explicit. (These were defined in the same way as in Study 2.) To assess the reliability of the coding, two independent raters coded one page (approximately 27% of the data) randomly selected from each subject's protocol. The percent of agreement was 83%, and differences were resolved through discussion. In general, there were no differences between experts and novices in the number of implicit propositions or explicit propositions produced. (Note that in Study 2, experts produced a greater number of implicit statements. The difference between the two results most likely resulted from the fact that the previous study examined the production of knowledge about familiar dinosaurs, whereas in this study, the dinosaurs were novel.)

Although there was no difference in the number of implicit propositions produced, there were notable differences in the content of the implicit statements. We categorized implicit propositions into those that were related to the domain of dinosaurs (e.g., "it eats meat"), and those that were related to world knowledge (e.g., "like a rhinoceros"). The majority of the experts' implicit propositions were domain-related (a mean of 19 out of 27), whereas the majority of the novices' implicit propositions were related to world knowledge (a mean of 26 out of 33), and the interaction between type of implicit proposition and knowledge was significant ($F(1,8) = 13.40, p < .01$; see Figure 5). Thus, the advantage of having domain knowledge is that experts can use it in making inferences and attributions about new dinosaurs, whereas novices must rely more on world knowledge.

Insert Figure 5 about here

In contrast with the Implicit propositions, the Explicit propositions referred to attributes of the dinosaurs that were directly observable from the picture. Experts and novices did not differ in the total number of Explicit propositions made (Exp. $M = 36.6$, Nov. $M = 20.4$, not significantly different by a Mann Whitney test), a result consistent with the results of Gobbo and Chi (1986). The consistency between the two studies in the explicit proposition results contrasts with the inconsistency in the implicit proposition results. This is to be expected, since explicit propositions do not rely on knowledge one already has about the dinosaurs, whereas implicit ones do. Thus, the introduction of novel dinosaurs in the third study had no effect on experts' and novices' production of explicit statements, but did erase the distinction between the number of implicit statements made by experts and novices. The experts, in effect, were not quite so "expert" when making inferences about the novel dinosaurs.

The Explicit statements could be differentiated as to whether they were domain related (e.g., "has sharp teeth") or world-knowledge-related (e.g., "has wrinkles"). The experts tended to produce more Explicit domain related propositions than novices, although this difference was not significant by a Mann Whitney test (Exp. $M = 17.8$, Nov. $M = 10.6$). We feel that this result is not statistically significant in part because we were very liberal in coding novices' responses. For instance, if they said "have teeth" as opposed to "have sharp teeth," they were credited with pointing out domain related features, even though it is not clear that they know the significance of "having sharp teeth" (something expert children would say). At any rate, the greater number of relevant features identified by the expert children suggests that they know which attributes to attend to, and they are less likely to mention irrelevant features like "is round in the middle".

In summary, we can draw two conclusions from the analyses of the propositions generated in the free generation task. First of all, there were no differences between experts and novices in the number of implicit and explicit propositions produced. The equivalence of these measures validates our assumption that we have succeeded in obtaining a homogeneous sample of children. For instance, neither group of children is more articulate than the other. Furthermore, there was no difference in the way expert and novice children used general world knowledge to constrain their inferences. That is, when expert and novice children did use general world knowledge to infer new information, they were equally proficient at making comparisons, and they both relied to the same extent on relevant animal-related knowledge to make attributions about the novel dinosaurs. However, the expert children excelled in one regard: they used predominantly domain knowledge (as opposed to general world knowledge) to constrain their inferences. Although we do not have longitudinal learning data, our current results imply that such domain-constrained inferences would allow expert children to learn about new domain concepts more quickly and accurately.

In contrast with the Implicit propositions, the Explicit propositions referred to attributes of the dinosaurs that were directly observable from the picture. Experts and novices did not differ in the total number of Explicit propositions made (Exp. $M = 36.6$, Nov. $M = 20.4$, not significantly different by a Mann Whitney test), a result consistent with the results of Gobbo and Chi (1986). The consistency between the two studies in the explicit proposition results contrasts with the inconsistency in the implicit proposition results. This is to be expected, since explicit propositions do not rely on knowledge one already has about the dinosaurs, whereas implicit ones do. Thus, the introduction of novel dinosaurs in the third study had no effect on experts' and novices' production of explicit statements, but did erase the distinction between the number of implicit statements made by experts and novices. The experts, in effect, were not quite so "expert" when making inferences about the novel dinosaurs.

The Explicit statements could be differentiated as to whether they were domain related (e.g., "has sharp teeth") or world-knowledge-related (e.g., "has wrinkles"). The experts tended to produce more Explicit domain related propositions than novices, although this difference was not significant by a Mann Whitney test (Exp. $M = 17.8$, Nov. $M = 10.6$). We feel that this result is not statistically significant in part because we were very liberal in coding novices' responses. For instance, if they said "have teeth" as opposed to "have sharp teeth," they were credited with pointing out domain related features, even though it is not clear that they know the significance of "having sharp teeth" (something expert children would say). At any rate, the greater number of relevant features identified by the expert children suggests that they know which attributes to attend to, and they are less likely to mention irrelevant features like "is round in the middle".

In summary, we can draw two conclusions from the analyses of the propositions generated in the free generation task. First of all, there were no differences between experts and novices in the number of implicit and explicit propositions produced. The equivalence of these measures validates our assumption that we have succeeded in obtaining a homogeneous sample of children. For instance, neither group of children is more articulate than the other. Furthermore, there was no difference in the way expert and novice children used general world knowledge to constrain their inferences. That is, when expert and novice children did use general world knowledge to infer new information, they were equally proficient at making comparisons, and they both relied to the same extent on relevant animal-related knowledge to make attributions about the novel dinosaurs. However, the expert children excelled in one regard: they used predominantly domain knowledge (as opposed to general world knowledge) to constrain their inferences. Although we do not have longitudinal learning data, our current results imply that such domain-constrained inferences would allow expert children to learn about new domain concepts more quickly and accurately.

Reasoning structures. We were interested in examining more directly how expert and novice children use their knowledge to infer information about the novel dinosaurs. To do this, we coded the free generation protocols at a higher level of analysis-- that of reasoning structures. Reasoning structures consisted of one or two sentences that together completed a reasoning chain. Reasoning chains were classified into four types: Family/Superordinate, Dinosaur Animal, and Attribute.

Family or Superordinate reasoning structures included all inferences that invoked a dinosaur family or superordinate class (such as meat-eating) from an attribute or a dinosaur, and vice versa. Hence, they included all of the following:

1. inferences from a superordinate class to an attribute (e.g., "He's pretty dangerous. [Why?] Cause he's a meat-eater.")
2. inferences from an attribute to a superordinate class (e.g., "I figure it's a planteater cause it needs weapons to protect from a meateater")
3. inferences from a superordinate class to a dinosaur (e.g., "It's not Triceratops. [Why] "Cause they're new dinosaurs")
4. inferences from a family to an attribute (e.g., "I guess he's probably a good swimmer" [Why?] "Cause duckbills are good swimmers")
5. inferences from an attribute to a family (e.g., "Duckbill" [Why?] "Cause it has this [beak]").

Dinosaur or Animal reasoning structures were inferences from an attribute to a dinosaur or to an animal, respectively (e.g., Dinosaur: "He's like Tyrannosaurus Rex. [Why?] Cause I saw his teeth."; Animal: "A lizzard. [Why?] Cause he has this long tail.").

Inferences from one attribute to another attribute were called Attribute reasoning structures. An example is: "He could walk real fast [Why?] Cause he has giant legs." The "giant legs" led directly to the "walking fast" inference, rather than proceeding via a Family/Superordinate or Dinosaur Animal inference.

There is an important difference between the first two types of inferences (Family/Superordinate, Dinosaur, and Animal) and the Attribute inferences. The difference is that for the first three types of inferences, the child is engaging in a type of *hierarchical* reasoning, whereas in the Attribute inferences only *linear* reasoning is involved. In hierarchical reasoning, a specific feature in the picture of the dinosaur triggers the identification of a familiar

animal, dinosaur, dinosaur family or superordinate class, and this identification then generates a new inference. For example, in the first quote listed above, the child said, "He's pretty dangerous. Cause he's a meateater." Some feature(s) in the dinosaur picture must have led the child to identify the dinosaur as a meateater, which in turn generated the inferred knowledge that this particular dinosaur must be dangerous, since meateaters in general are dangerous. Similarly, for the Dinosaur inferences, to say that "He's like Tyrannosaurus Rex. Cause I saw his teeth" implies that the "teeth" in the picture allowed the child to identify the novel dinosaur as similar to Tyrannosaurus Rex. Presumably, had we probed further, the child might then have made other attributions based on the inference that the dinosaur was like Tyrannosaurus.

Attribute inferences, in contrast, do not involve an initial categorical or instance identification. They are *linear* in nature, rather than hierarchical. To say that a dinosaur "Could walk real fast cause he has giant legs" implies that the explicit feature "giant legs" leads directly to the inference that the dinosaur "walks fast". This inference is not hierarchical because the child did not identify "giant legs" as an attribute of a certain class of dinosaurs, and then make the inference that because this class of dinosaurs walks fast, this particular dinosaur walks fast.

Analysis of the reasoning chain at this sentence level revealed that experts and novices made approximately the same number of inferences (Exp. $M = 9.2$, Nov. $M = 11.4$, $U = 12.0$, $p > .1$). Consistent with the propositional analysis, this level of analysis again shows that both experts and novices were equally capable of generating information about the new dinosaurs using knowledge they had. In particular, experts and novices did not differ in the number of times they reasoned analogically to animals (Exp. $M = 0.8$, Nov. $M = 0.6$). Thus, this validates our assumption that the two groups have equivalent amounts of knowledge about animals in general, and that they were equally proficient at reasoning from animal instance knowledge.

However, because the experts had more knowledge about dinosaurs, they were able to make use of this knowledge by generating a significantly greater number of Dinosaur reasoning chains than the novices (Exp. $M = 3.6$, Nov. $M = 0.2$, Mann Whitney $U = 3.5$, $p < .05$, one-tailed). In fact this was their largest category, accounting for 39.2% of the experts' inferences. Experts also made a larger number of Superordinate/Family inferences than the novices, although the difference was not significant (Exp. $M = 2.2$, Nov. $M = 1.0$). We suspect that a significant difference would have been reached had we used more novel dinosaurs that were prototypical. At any rate, if we combine both the Dinosaur inferences and the Family/Superordinate inferences, the expert children made a significantly greater number of these domain-related hierarchical type inferences than the novices. (Exp. $M = 5.8$, Nov. $M = 1.2$, $U = 2.5$, $p < .05$, one tailed). In contrast, novices made many more Attribute inferences than experts

(Exp. \underline{M} = 2.6, Nov. \underline{M} = 9.6, \underline{U} = 4.0, $p < .05$, one-tailed). The Attribute inference category was clearly the largest for the Novices, accounting for 78.6% of their inferences.

Thus, the main finding from this analysis is that experts' reasoning structures are derived from categorizing the novel dinosaur into a domain-related type or a class, whereas novices' inferences are not constrained by such instance or family knowledge. What this means is that expert children, in examining a picture of a novel dinosaur, seek out "critical" features that can enable them to identify it as similar to a particular dinosaur instance or member of a family. Additional inferences are then generated on the basis of this initial classification. Because novice children lack the domain related instance or family knowledge, they cannot seek specific features that can lead to classification of the dinosaur. This forces the novices to generate inferences on the basis of individual attributes. Their inferences are necessarily more limited in scope, because, for example, there are only a few things one can say about a dinosaur having "giant legs". There is no class of dinosaurs having giant legs, which could lead to a large number of inferences. Thus, the use of Attribute reasoning necessarily limits the potential inferences novices can generate about a novel dinosaur.

Causal statements. Another way to analyze children's reasoning in the free generation task is to look at their causal explanations. We classified as causal statements all statements that were spontaneous (i.e., not prompted by an experimenter) and that were connected by the words "because", "cause", "so", and "since". A typical causal statement identifies a dinosaur feature or places the dinosaur into a category in the premise, then follows with an explanation for the identification, such as "It's a meateater because it has sharp teeth."

In general, there was no difference between the number of causal statements generated by the expert and novice children (Exp. \underline{M} = 2.8, Nov. \underline{M} = 6.6; not significant by a Mann Whitney test). (The greater mean for the novices is inflated by one subject's data.) Causal statements were classified into three categories: (a) Domain related premise and domain related explanation (Dom Rel-Dom Rel: e.g., "It's a meateater because it has sharp teeth."); (b) Domain related premise and explanation drawn from world knowledge (Dom Rel- WK: e.g., "It's a meateater because it has 14 toes."); and (c) World knowledge premise and world knowledge explanation (WK-WK: e.g., "It lives in the desert because it looks like a lizzard."). There was no difference between experts and novices in the number of WK-WK causal statements (Exp. \underline{M} = 0.2, Nov. \underline{M} = 3.0, not significant by a Mann Whitney test). For those causal statements (categories a and b) in which the children identified a domain-related feature or classification in the premise, there was a difference in the kinds of explanations novices and experts gave. Experts gave explanations based predominantly on domain-related knowledge (category a), whereas novices gave explanations based predominantly on world knowledge (category b) (See Figure 6). The interaction was significant

(Exp. means: Dom Rel- Dom Rel = 2.2, Dom Rel- WK = 1.0; Nov. means: Dom Rel-Dom Rel = 0.4, Dom Rel - WK = 2.6 ; $F(1,8) = 7.71$, $p = .024$.) Hence, even though both novices and experts are equally facile at producing causal inferences, the explanations given by novices are less appropriate for justifying the identified feature. The use of world knowledge necessarily limits the accuracy of the novices' causal explanations.

Insert Figure 6 about here

Accuracy within dinosaur taxonomy. Most of our analyses have focused on the extent to which experts' inductions are hierarchical in nature and are based on domain-related knowledge. It is important to know whether experts' greater domain knowledge leads to inferences that are more accurate than novices' with respect to an absolute scientific standard. Consequently, we examined the accuracy of children's generated knowledge in the free generation protocols. For example, a statement such as "Usually dinosaurs live on land if they have small heads like that" is clearly incorrect because some dinosaurs with small heads live in the water. A statement such as "[It eats] plants because it's low down and plants are low down" is not a scientifically acceptable reason for classifying a dinosaur as planteating. Similarly, the statement, "He'd be able to eat animals, might eat plants too" is not scientifically acceptable because no dinosaur eats both animals and plants.

The experts did not produce *any* incorrect statements in the free generation task. Each novice produced an average of 1.8 incorrect statements (Mann-Whitney $U = 5.00$, $p < .05$). This accuracy analysis indicates that experts have a stable body of domain knowledge that allows them to generate correct inferences about novel stimuli that are consistent with what they know. Novices, even though they actively attempt to generate inferences and generalizations, produce incorrect inferences because they lack the relevant domain knowledge to constrain them.

Results from the Sorting Tasks

Sorting of familiar dinosaurs. We speculated from the first study that children, as they acquire expertise, will structure their dinosaur knowledge hierarchically. Individual dinosaurs are grouped into families, and the families are then coalesced into higher categories of meat- and plant-eaters (a differentiation zoologists also use to classify mammals). Across the second and third studies, we have data on 20 children's sorting of the 20 better-known dinosaurs. Ten of the children were novices and 10 were experts, and they ranged in age from 4 years to 7 years.

To gain a rough idea of what hierarchical level (if any) children's sorting patterns fit into, we first classified each child's sorting into Superordinate, Emerging Superordinate, Family, Emerging Family, or No Classification. To be classified as Superordinate, a child's groups had to satisfy the following three criteria:

1. Half or more of the members of each group must be from the same superordinate category.
2. Half or more of the members of each superordinate category must be grouped together.
3. Members of each superordinate category must form a majority (more than half) of at least one group of two or more members.

The same criteria were used for classification as Family, substituting "family" for "superordinate." If a child's sorting satisfied 2 out of the 3 criteria for Superordinate and consisted of 2 or 3 groups, it was classified as Emerging Superordinate. Similarly, if a child's sorting satisfied 2 out of the 3 criteria for Family and consisted of 4 or more groups, it was classified as Emerging Family. Sortings that did not conform to these criteria were designated No Classification, and were not considered to fit a hierarchical scheme.

Based on these criteria, 3 of the 10 experts classified the 20 dinosaurs into the Superordinate meat-eating and plant-eating categories. This replicates the sorting result found for the expert in the first study. None of the novices spontaneously grouped the dinosaurs on the abstract dimension of diet, and only one novice had an Emerging Superordinate grouping. When asked to subdivide their Superordinate groupings, the three experts created Family or Emerging Family groupings. We know, then, that at least some experts can group the dinosaurs at different hierarchical levels when requested. Novices do not appear to have similar abilities.

Four experts initially grouped the dinosaurs by families and an additional one had an Emerging Family grouping. Only one novice's grouping could be classified as Family, but four novices showed Emerging Family distinctions. It is not surprising that novices have an emerging ability to sort dinosaurs on the basis of family membership, since family membership is determined to a large extent by visual similarity. Dinosaurs belonging to the same family tend to have similar visual features, such as duck-like bills for the Duckbill family, and notably distinct horns for the Horned dinosaurs. Wattenmaker, Nakamura and Medin (in press) also believe that one can use apparent family resemblance to sort objects into categories. In some cases, however, (as in the case of our expert children) the presence of external visual resemblance may have masked the use of deeper principles, so that it appears as if the novices and experts were using the same principles to guide their sortings.

If we combine the results, we find that experts sorted strictly by family or superordinate category, whereas only one novice made strict classifications. In contrast, five novices showed emerging family or emerging superordinate distinctions, whereas only one expert did so. The difference was significant by a Fisher's Exact test, p

< .05. This result suggests that novices were forming groups on-line, as the task of sorting demanded it of them. In contrast, more of the expert children appeared to sort on the basis of some predefined representation of family or superordinate structure.

A more detailed quantitative analysis, performed only on the data from the 5 novices and 5 experts from the third study, compared the number of "perfect" groups formed by experts and novices. The analysis was performed on the groupings children constructed after they were asked to break down an initial large group. A perfect group is defined as one that includes all of the dinosaurs from one family and no dinosaurs from other families. Experts produced significantly more perfect groups than novices (Experts $\bar{M} = 2.2$; Novices $\bar{M} = 0.6$, Mann Whitney $U = 3.00$, $p < .05$, one-tailed). Conversely, the novices had three times as many groups with intrusions (groups with a mixture of dinosaurs from different families) as experts (a mean of 0.8 mixed groups for experts and 2.4 mixed groups for novices, $U = 4.00$, $p < .05$, one-tailed). This finding implies that the expert children are less likely to make "errors" in their sorting, which would presumably occur if one was relying strictly on visual similarity as a basis for the groups. We interpret this as further evidence that family structures that already existed in the experts' knowledge base guided their sorting, whereas the novices had to rely on properties they observed on-line to come up with family-like groupings.

In order to capture the differences in the principles which guide the sortings of expert and novice children, we also examined the explanations for their sortings they gave for their groupings. Even though both experts and novices appeared to sort on the basis of family resemblance, the novices relied more on properties of the dinosaurs visible in the pictures. They explained by saying:

"They all look alike. Their heads look alike."

"They have small hands...And their skin is rough."

"They have the same heads...And they have horns."

Expert children, on the other hand, explained their groupings on the basis of more abstract dimensions that were implicit rather than explicitly visible in the picture:

"This group...is protected and harmful dinosaurs that are planteaters."

"These dinosaurs are plant-eaters. They aren't harmful."

"These are the only meat-eaters I could find."

This particular expert child's protocol not only shows that the child attended to more implicit, abstract features, but it also shows some hierarchical embedding. That is, planteaters can be harmful or not harmful.

The difference between experts' and novices' explanations for their sortings held up across both the second and third studies. In the second study, novices gave explicit physical features 83.2% of the time, whereas experts did so only 42.6% of the time ($F(1,8) = 7.176, p < .05$). Experts mentioned implicit features such as diet, defense, and habitat more often than novices (Exp. $M = 77.6\%$, Nov. $M = 16.8\%$, $F(1,8) = 21.605, p < .01$). In the third study, 40% of the expert children's reasons concerned implicit features as compared to 13% of the novices' reasons ($F(1,8) = 3.495, p < .05$, one-tailed).

Thus, the analyses of the explanations further support our interpretation that the novices are sorting on-line, basing their groupings on the visible explicit features, in contrast to the experts' more principled sorting based on implicit features. In conclusion, the analyses of the sorting of the *familiar* dinosaur data suggest the following. It is possible for novices to sort objects on the basis of family resemblance, without knowing or understanding the underlying principle for the categorization. This reflects an ability to perceive similarity in the visual stimuli, and is a critical attribute for humans to possess in order to learn categories for objects. Expert children, however, have additional knowledge beyond the explicit visual features of the dinosaurs. They sort on the basis of their internal representation of the family structures. In addition, a few expert children sort according to diet, an even higher level of classification that coalesces family structures, in the same way that zoologists for classify animals (Storm, 1978).

Forced choice sort. The effect of having a coherent representation of family structures can also be detected by examining the way children sort unfamiliar dinosaurs. In sorting the unfamiliar dinosaurs, the children had the choice of either inserting the unfamiliar ones into the existing categories created by the Familiar dinosaur sorting, creating new categories, or partially dissolving the existing categories and forming revised ones. The end result was that both groups of children formed a total of 9-10 groups. Although there were no quantitative differences between the expert and novice children, there were some differences in how accurately they sorted the unfamiliar dinosaurs, especially the prototypical ones. One measure of accuracy is the degree of overlap among the features of each group's members. A rough measure of the overlap of features of the members of a group is to measure the degree to which members match on the following four attributes: Body, Head/Teeth, Legs, and Other (all other features, including armor, horn, neck, tail, crest, etc.) Each of the four attributes of each dinosaur was given a family classification. For example, one of the conflict dinosaurs had a body of a Giant Planteater, had the head/teeth of a Horned dinosaur, had Meateater legs, and had other features like a Giant Planteater. To determine the degree of overlap for a child's grouping, each attribute for each dinosaur in a group was given a score representing the

percentage of dinosaurs in the group having that particular family classification for that attribute. Then all the scores for a group (four for each dinosaur) were averaged to come up with a group mean. Finally, the group means were averaged for each child to come up with an overall percentage overlap. The five *newly-discovered* dinosaurs were excluded from the analysis. Groups with only one member were also excluded.

For the expert children, the percentage of overlap was 78.5%, whereas for the novice children, the overlap was 61.4% (Mann Whitney $U = 2.0$, $p < .05$). The significant difference between experts and novices in the percentage of overlap suggests that even though category membership can be determined to a large extent by the visual similarity among the dinosaurs, the expert children were still superior at detecting such similarity. Experts were probably better at such detection because they imposed their schemata of family structures to help them encode the relevant features. That is, they based their groupings of the dinosaurs on features that are predetermined by their family schemata. Novice children have to rely more on "bottom-up" processing of the physical features to make decisions about category membership, and as a result may miss or underemphasize some of the relevant features.

Compared to novices, experts were especially good at classifying the unfamiliar *prototypical* dinosaurs. Experts placed an average of 36% of the new prototypical dinosaurs with a group which included all the familiar dinosaurs from the same family as the prototypical dinosaur, without any intrusions of familiar dinosaurs from other families. (Sometimes children placed some other unfamiliar dinosaurs with the group as well.) Novices, in contrast, did not insert any of the new prototypical dinosaurs in such groups. The difference was significant by a Mann Whitney test, $U = 0$, $p < .01$. Thus, experts were able to use their knowledge of families to infer the correct family grouping for dinosaurs they were not familiar with.

Experts clearly distinguished between the *newly-discovered* dinosaurs and the *prototypical* dinosaurs. They classified 80% of the *newly-discovered* dinosaurs into new groupings which were not formed during the sorting of the Familiar dinosaurs. These new groups did not include any familiar dinosaurs. In contrast, they classified only 20% of the *prototypical* dinosaurs into such new groupings ($p = .031$ by a Wilcoxon's test). Novices, on the other hand, did not differ in their placement of the *newly-discovered* and *prototypical* dinosaurs into the new groupings (Newly-discovered: 80%, Prototypical: 52%, difference not significant by a Wilcoxon's test).

In sum, experts were able to utilize their family schemata to sort dinosaurs they had never seen before. They not only were better at detecting similarities between the dinosaurs, but they placed prototypical dinosaurs into the correct existing family groupings, and formed new groups for the newly-discovered ones. Such discrimination between the *prototypical* and *newly-discovered* dinosaurs was not evident in novices' sorting. This suggests that the

experts would be more likely to use categorical reasoning to constrain their inferences for the *prototypical* than for the *newly discovered* dinosaurs.

Analyses of Protocols Across All Tasks

In this section, we analyze the protocols generated by the children across all the tasks. That is, the remarks and explanations given in the Free Generation Task, the Probe Task, the Sorting tasks, and the Oddity task, are all combined for most of the analyses to be discussed below. The variety of analyses provide additional support for expert knowledge of family and superordinate diet classification.

Inferences based on families. Of the 20 novel dinosaurs, 5 prototypical and 7 nonprototypical ones had features that were all consistent with either the meateating category or the planteating category. Across the free generation and probe tasks, experts were more correct than novices in classifying the 5 prototypical dinosaurs according to diet. (Exp. \bar{M} = 95% correct, Nov. \bar{M} = 44.6% correct, $U = 1.0$, $p = .013$). When determining the diet class of the 7 nonprototypical dinosaurs, experts and novices were equally correct (Exp \bar{M} = 70.0% correct, Nov. \bar{M} = 70.6% correct). The fact that experts have an advantage for prototypical dinosaurs but not for nonprototypical dinosaurs suggests that experts cannot be relying strictly on individual attributes (such as "big head") to make diet classifications, since individual attributes can be used equally effectively for prototypical and nonprototypical dinosaurs. For the prototypical dinosaurs, experts may be able to draw upon their knowledge of families to categorize each dinosaur and thereby identify its diet. This result is consistent with the previous analysis showing that experts sort prototypical dinosaurs better than novices. The finding is also consistent with our previous conclusion that expert children use reasoning structures that are hierarchical, whereas novices use linear reasoning, in which inferences are based on individual attributes.

Additional evidence that the experts in Study 3 were using family-level knowledge at least part of the time to infer the dinosaurs' diet can be gleaned from the reasons they gave for their diet classifications in the free generation and probe tasks. The expert children used membership in a family to justify diet classifications 14.4% of the time, whereas novices never did (Mann-Whitney $U = 2.5$, $p < .05$). For example, an expert might say, "He eats plants because he looks like a Duckbill."

We also looked more generally at the total number of family-related statements children produced (e.g., "He's a Horned dinosaur.") Tabulating statements generated across all the tasks of this study, the expert children produced an average of 13.6 family related statements, whereas the novices produced an average of only 1.0 (Mann-Whitney $U = 5.0$, $p < .05$). Again, this suggests that experts more often make inferences about a novel dinosaur on the basis of its relation to a certain family.

Presence of superordinate diet classes. As in Study 2, we found that experts had a sense of the contrastive nature of the two superordinate diet classes, meat-eating and plant-eating, whereas novices did not. Across all the tasks, all five experts used the presence of a feature as justification for placement of a dinosaur in one diet class, and the absence of the feature as justification for placement of the dinosaur in the complementary diet class. Only one novice used features in this contrastive way ($p < .05$ by a Fisher Exact Probability test).

We can also detect the presence of superordinate meat- and plant-eating categories in the Oddity task. This task included 4 trials on which one dinosaur could be segregated from the other three on the basis of diet category. Experts grouped meat and plant-eaters in separate groups on 85% of these trials, while novices grouped meat and plant-eaters separately on only 55% of the trials (Mann-Whitney $U = 1.5$, $p < .05$).

Summary

In the third study of this series, we focused on how expert and novice children infer new domain-related information, by asking children to generate attributions about novel dinosaurs. The expert and novice children were selected so that they were matched on a number of dimensions indicating that they had equivalent abilities to learn. As the result of obtaining a matched sample, we detected no differences between the experts and the novices in a number of quantitative measures, such as the total number of propositions produced, the number of groups created when sorting dinosaurs, and the number of causal statements made. Nor did we detect any differences in the use of general learning strategies when domain knowledge was not involved, such as the use of comparisons, to animals, and the use of relevant features of world knowledge to constrain inferences.

However, experts showed a great advantage over novices in using domain knowledge to generate inferences. In particular, the expert children used domain-related knowledge to generate comparisons whereas novice children use world knowledge. Experts based their explanations of causality on relevant domain features to a greater extent than novices. The expert children knew what explicit features to attend to in the pictorial representations, and were able to use categorical reasoning to induce attributes about novel dinosaurs. Expert children based their sorting on well-defined schemata of family types, whereas the novice children tended to sort on-line, on the basis of visual similarity. Experts also made more family-related statements across all the tasks than novices. The experts' use of existing family schemata to guide sorting can be gleaned from (a) the existence of a few expert children's sort on an abstract (non-visual) diet classes, (b) the way expert children formed "perfect" groupings without intrusions, (c) the greater degree of overlap in the features among members of a group, and (d) the fact that their groupings of the novel dinosaurs conformed to their existing schemata of the family structures. That is, novel *prototypical* dinosaurs tended to be sorted with existing groups from the same family, whereas *newly-discovered* ones tended to form new groups, since no schemata presumably exist for these newly-discovered ones yet. These qualitative differences in

the use of existing knowledge clearly confer an advantage to expert children in learning new domain-related knowledge. Experts can learn new domain-related concepts more efficiently and accurately than novices. Thus, the third study in this series has shown how background knowledge per se can enable the expert children to learn new domain-related concepts more readily, despite the fact that both expert and novice children have the same fundamental learning skills.

Discussion

In a series of three studies, we have chosen to examine a very small, simple, and confined domain--dinosaur concepts, in order to assess how it may be structured in children's representations, and how the knowledge may be used. Even for such a simple domain, in which the knowledge is basically organized in a classification scheme, we found significant differences in the way experts and novices' representations of such knowledge are structured. Across the three studies, we have evidence to suggest that the structure of the expert children's knowledge is more coherent, both locally and globally (hierarchically). Local coherence can be assessed in a number of ways: by the number of links between concepts ... in the same family and the sharing of features that are common within the same family, by the use of connecting words in the production of discourse about a familiar dinosaur, and in the frequency with which the discussion of one attribute (such as defense) leads to the discussion of another attribute. Local coherence can further be assessed by the degree to which family structures are well-defined. This can be seen in the amount of family overlap of features of members of the same group, and the extent to which the groups conform to perfect family structures.

Global coherence or hierarchy can be assessed by the presence of diet-level classification schemes as well as family-level. Some experts were able to sort dinosaurs at both the superordinate and family levels, whereas no novices did this. Expert children could reason categorically, using superordinate, family, and dinosaur categories. Expert children used a single feature set of features to contrast Global coherence can also be seen in the way the superordinate meateating and planteating classes.

By working with a matched sample of children, we were further able to show that the novice children were just as competent as the expert children in using general learning skills, such as making comparisons, giving causal explanations, when these skills were assessed in a domain and (such as in the animal domain) in which both the expert and novice children had equivalent knowledge. The use of learning skills seems to differentiate the experts from the novices only when they interact with dinosaur knowledge. For instance, categorical reasoning can only be manifested if children have the categories of family types represented in memory, otherwise they would have to resort to linear reasoning. Likewise, sorting based on family types can only be manifested if these family types exist

in memory, otherwise sorting has to be based on visual similarities that can be detected on-line. Using a single set of features to contrast two diet classes can only be done if one has a representation contrasting the two classes. (That is, only if the classes are locally coherent, in that members of the same class share many similar features, and members of contrasting classes share few features).

In sum, we believe that we have marshalled sufficient evidence from the present set of three studies to suggest that children can act more or less intelligently on the basis of what knowledge they have. Even for such a small and confined domain of knowledge, a modest mastery of the domain is manifested in more intelligent explanations, constrained inferences, categorical reasoning, hierarchical classification, and classification based on well-defined family structures.

Novices' explanations, on the other hand, seem inconsequential, and their inferences are not constrained; therefore appear incorrect and irrelevant. Their linear reasoning is limited in scope, their classification is erratic, that it contains a number of intrusions, and they have no awareness that newly-discovered dinosaurs should be discriminated from prototypical ones.

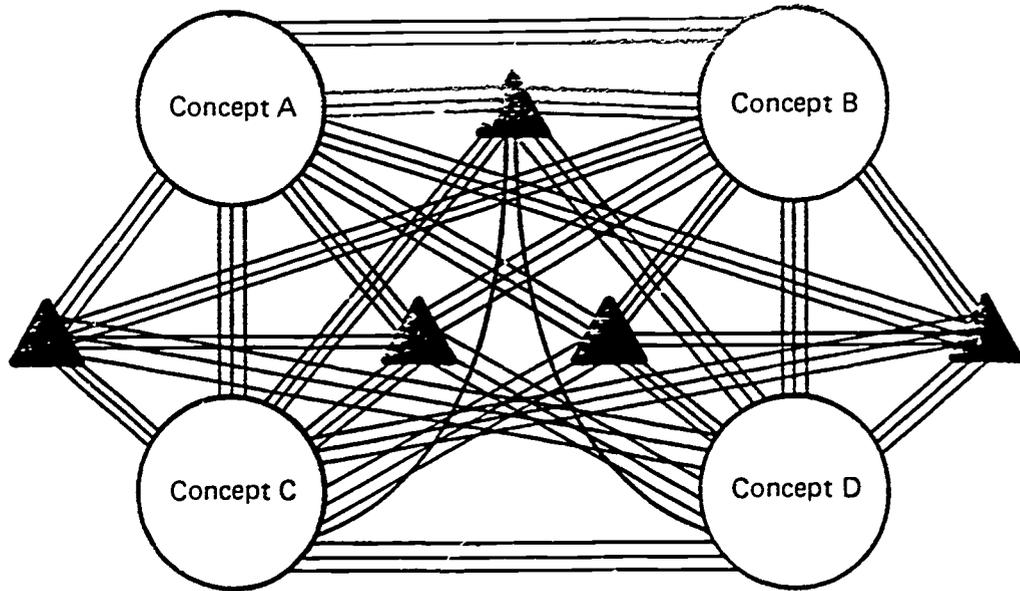
Because we have carefully controlled our subject samples so that they differ mainly in the amount of domain knowledge, a number of quantitative measures do not exhibit differences. Hence, the differences that we do see reflect the presence and structure of domain knowledge. The conclusion that one can draw is that one reason that children generally manifest global inadequacy across a number of domains is that they lack the relevant knowledge in a number of domains. By selecting a domain that some children know something about, we have demonstrated qualitatively superior abilities that can be attributed only to domain-specific restructuring.

References

- Bousfield, W. A. (1953). The occurrences of clustering in the recall of randomly arranged associates. *Journal of General Psychology*, 49, 229-240.
- Chi, M. T. H. (1983). Knowledge-derived categorization in young children. In D. R. Rogers & J. A. Sloboda (Eds.), *The acquisition of symbolic skills* (pp. 327-334). New York: Plenum Press.
- Chi, M. T. H. (in press). Children's lack of access and knowledge reorganization: An example from the concept of animism. To appear in F. Weinert & M. Perlmutter (Eds.), *Memory development: Universal changes and individual differences*. Hillsdale, NJ: Erlbaum.
- Chi, M. T. H., & Ceci, S. J. (1987). Content knowledge: It's representation and restructuring in memory development. In H. W. Reese & P. Lipsett (Eds.), *Advances in Child Development and Behavior*, (Vol. 20) (pp. 91-142). New York: Academic Press.
- Chi, M. T. H., Glaser, P., & Farr, M. (in press). (Eds.), *The nature of expertise*. Hillsdale, NJ: Erlbaum.
- Chi, M. T. H., & Koeske, R. (1983). Network representation of a child's dinosaur knowledge. *Developmental Psychology*, 19, 29-39.
- Gelman, R., & Baillargeon, R. (1983). A review of some Piagetian concepts. In J. H. Flavell & E. M. Markman (Eds.), *Handbook of child psychology: (Vol. III.) Cognitive development* (pp. 167-230). New York: Wiley.

- Gobbo, C., & Chi, M. T. H. (1986). How knowledge is structured and used by expert and novice children. *Cognitive Development, 1*, 221-237.
- Storm, C. (1978). Acquiring principles of semantic organization. *Journal of Experimental Child Psychology, 25*, 208-223.
- Wattenmaker, W. D., Nakamura, G. V., & Medin, D. C. (in press). Relationships between similarity-based and explanation-based categorization. In D. Hilton (Ed.), *Science and natural explanation: Commonsense conceptions of causality*. London: Harvester Press.

STRUCTURE I A



STRUCTURE II A

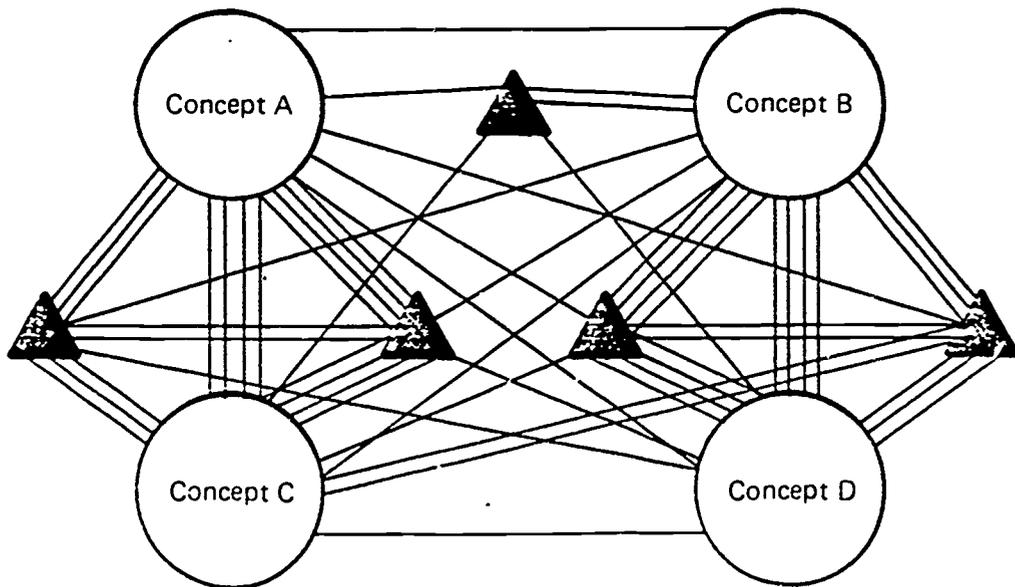
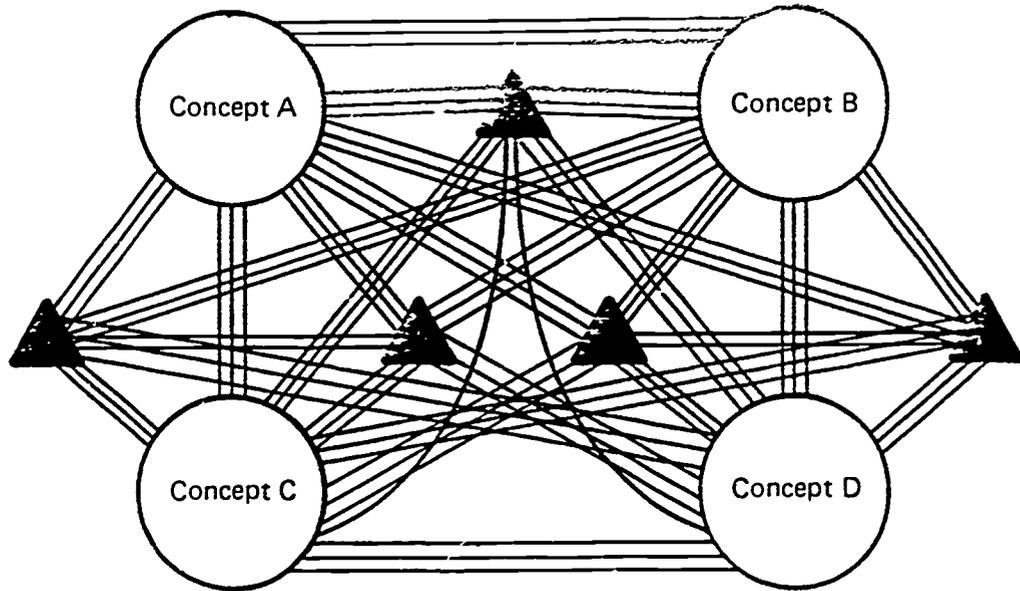


Figure 2

STRUCTURE I A



STRUCTURE II A

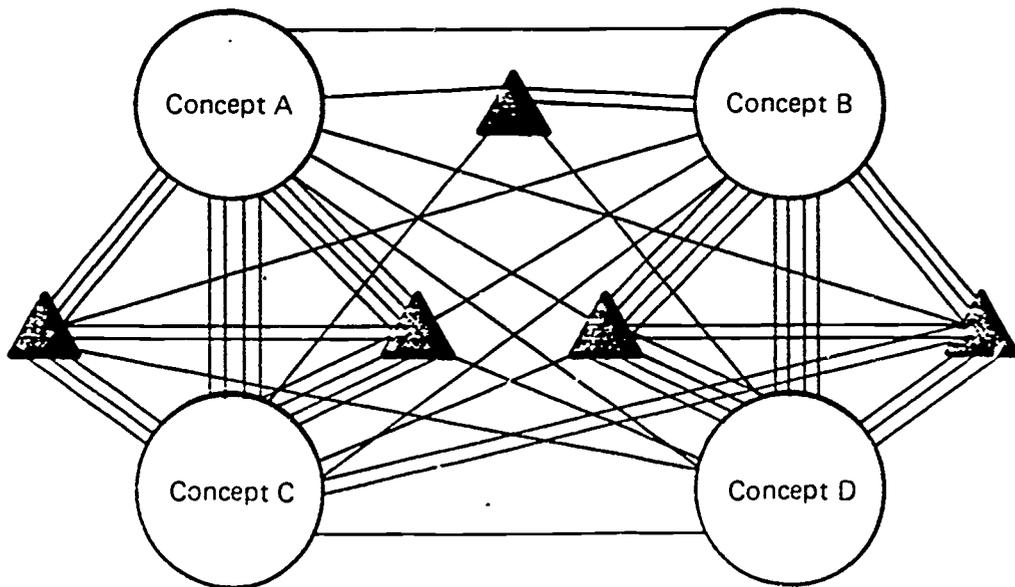


Figure 2

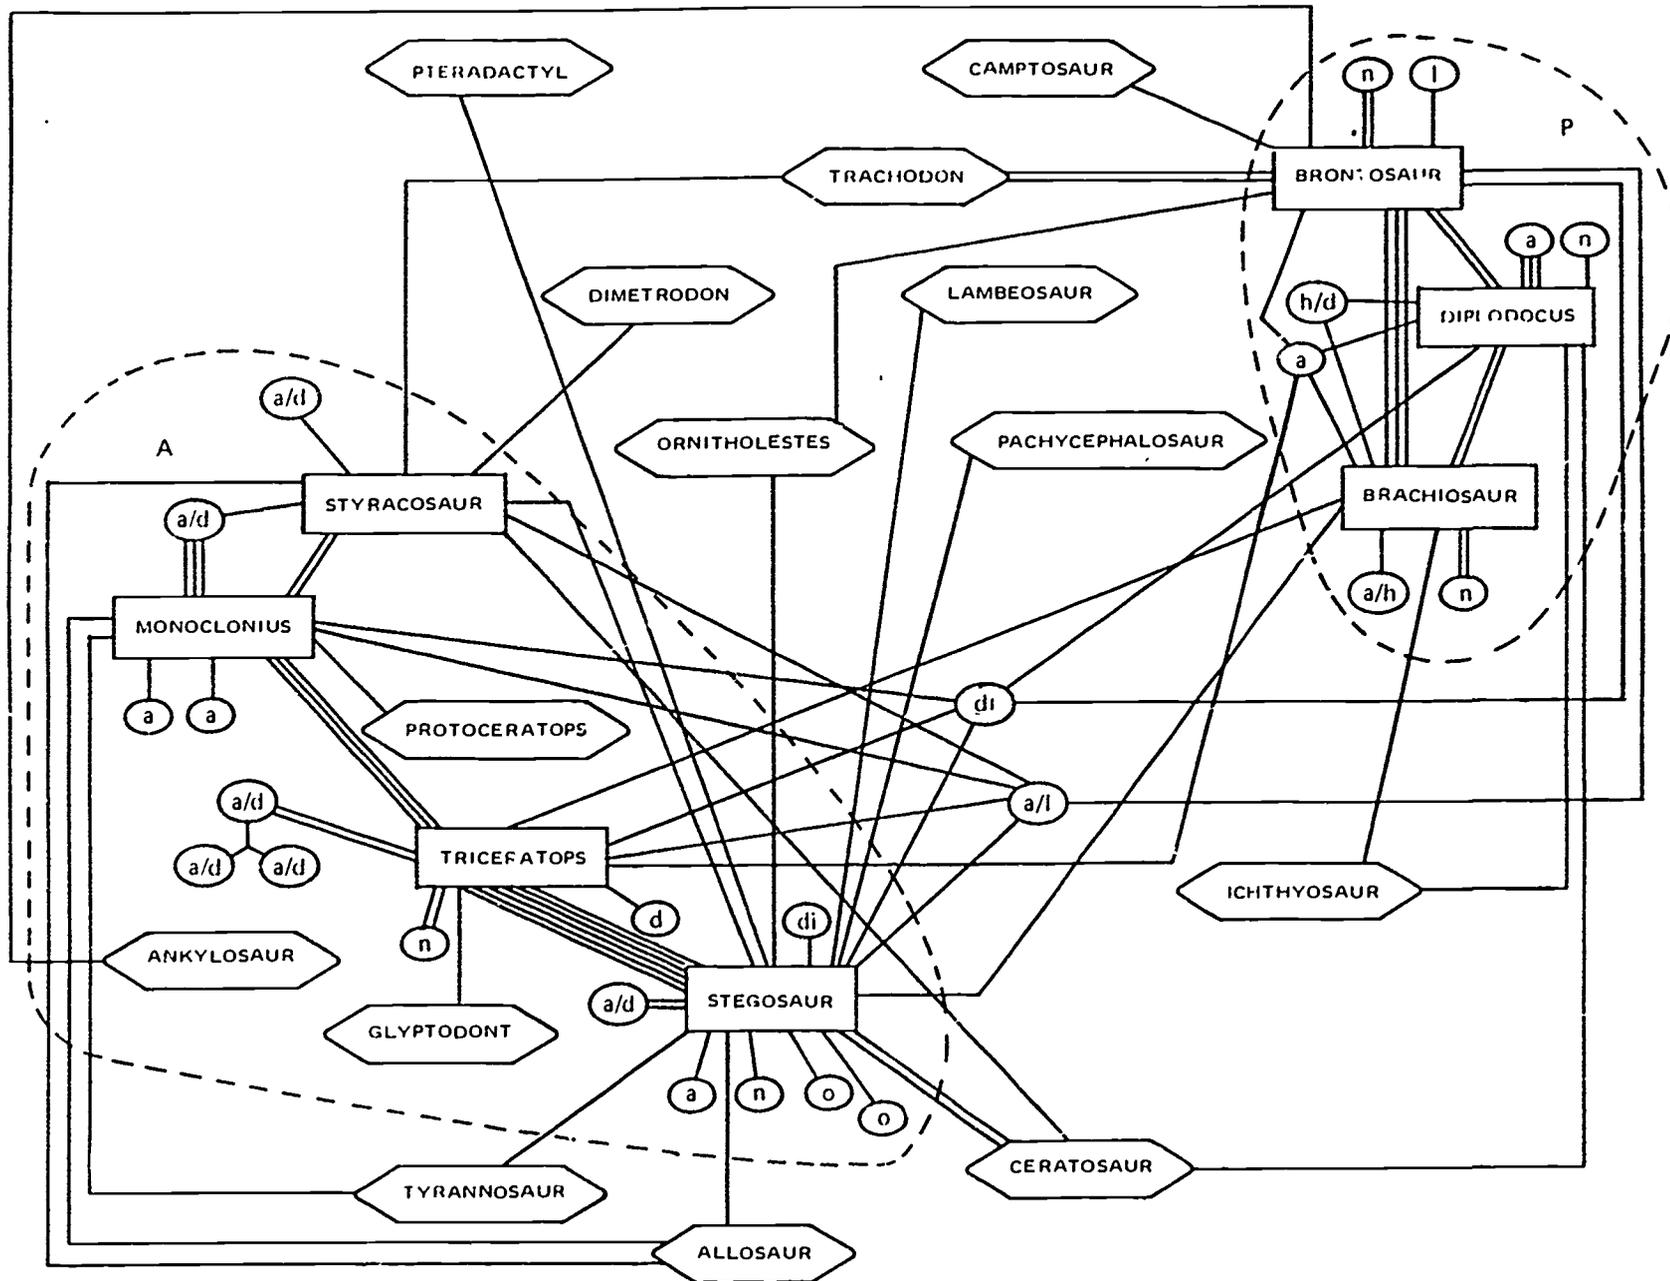


Figure 3

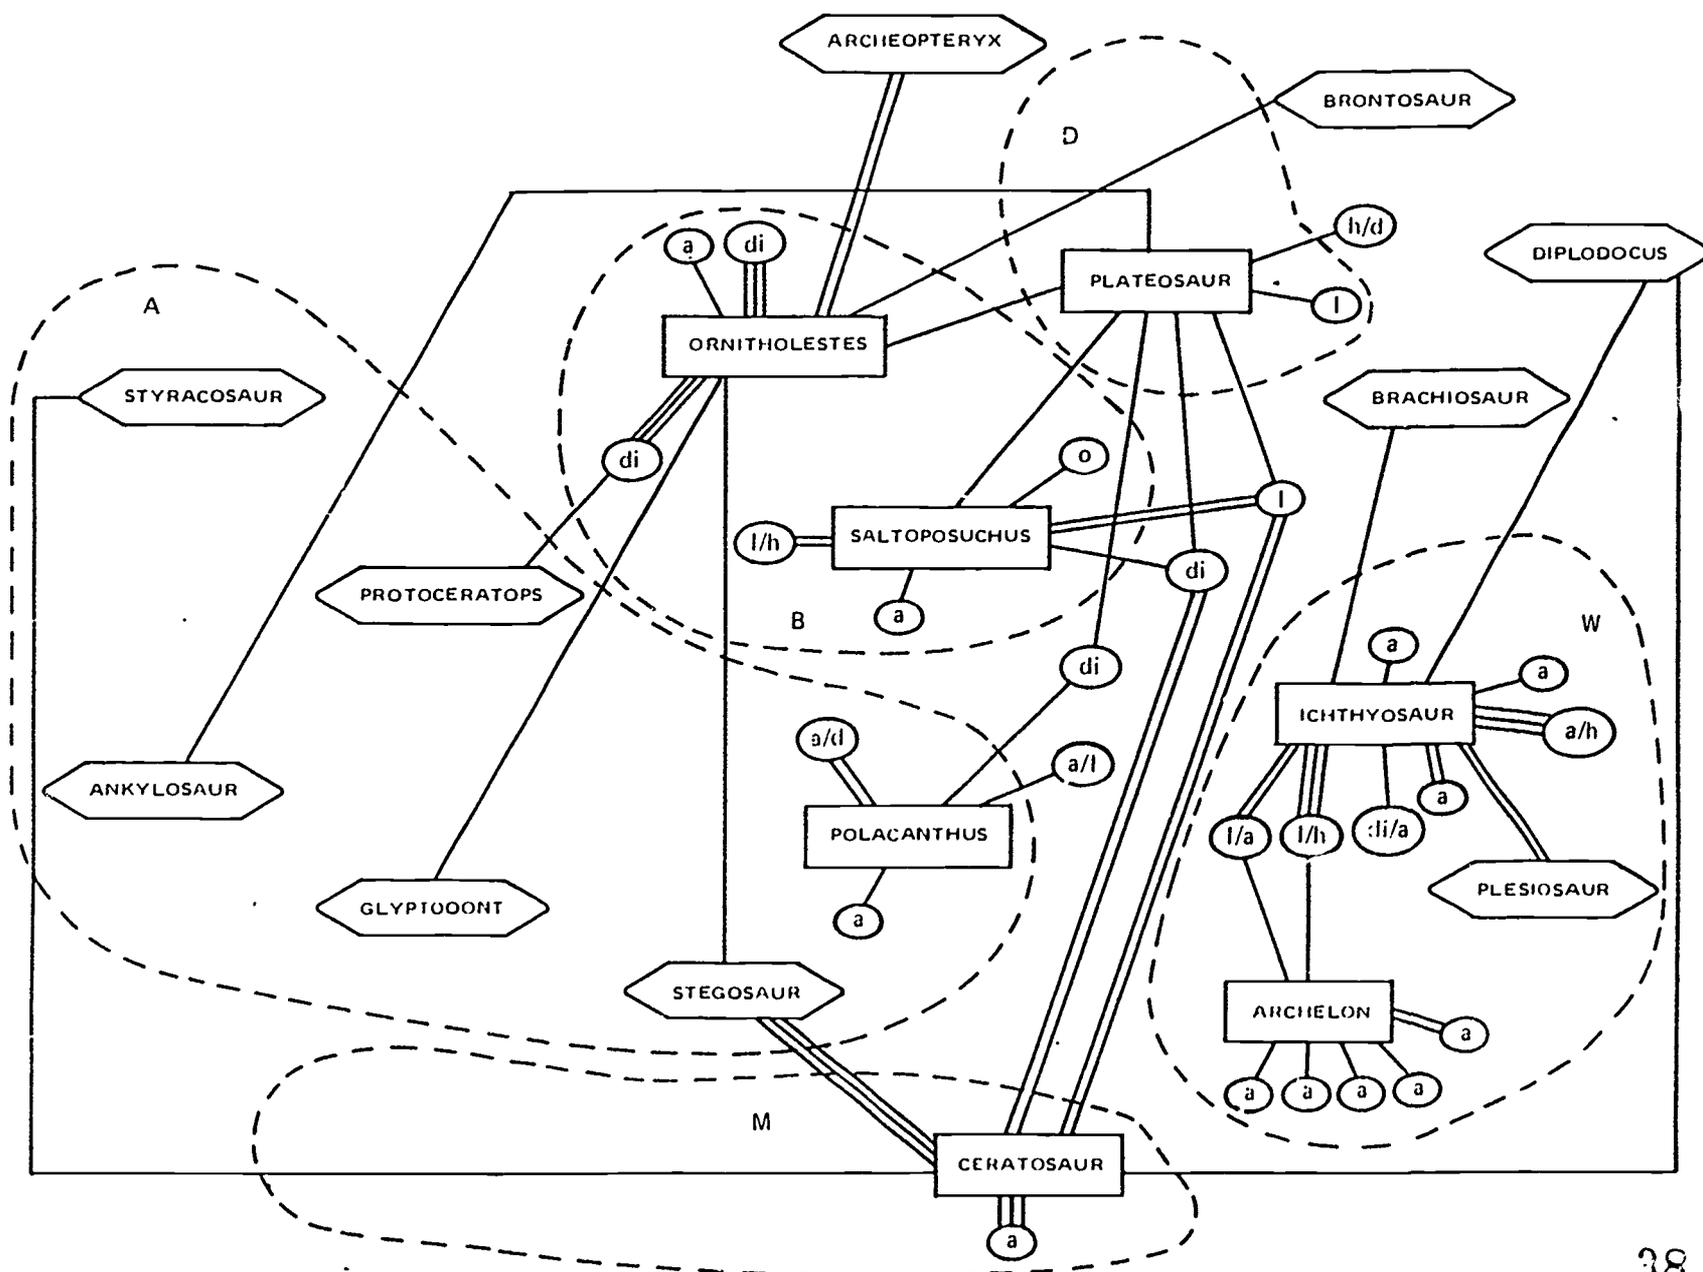


Figure 4

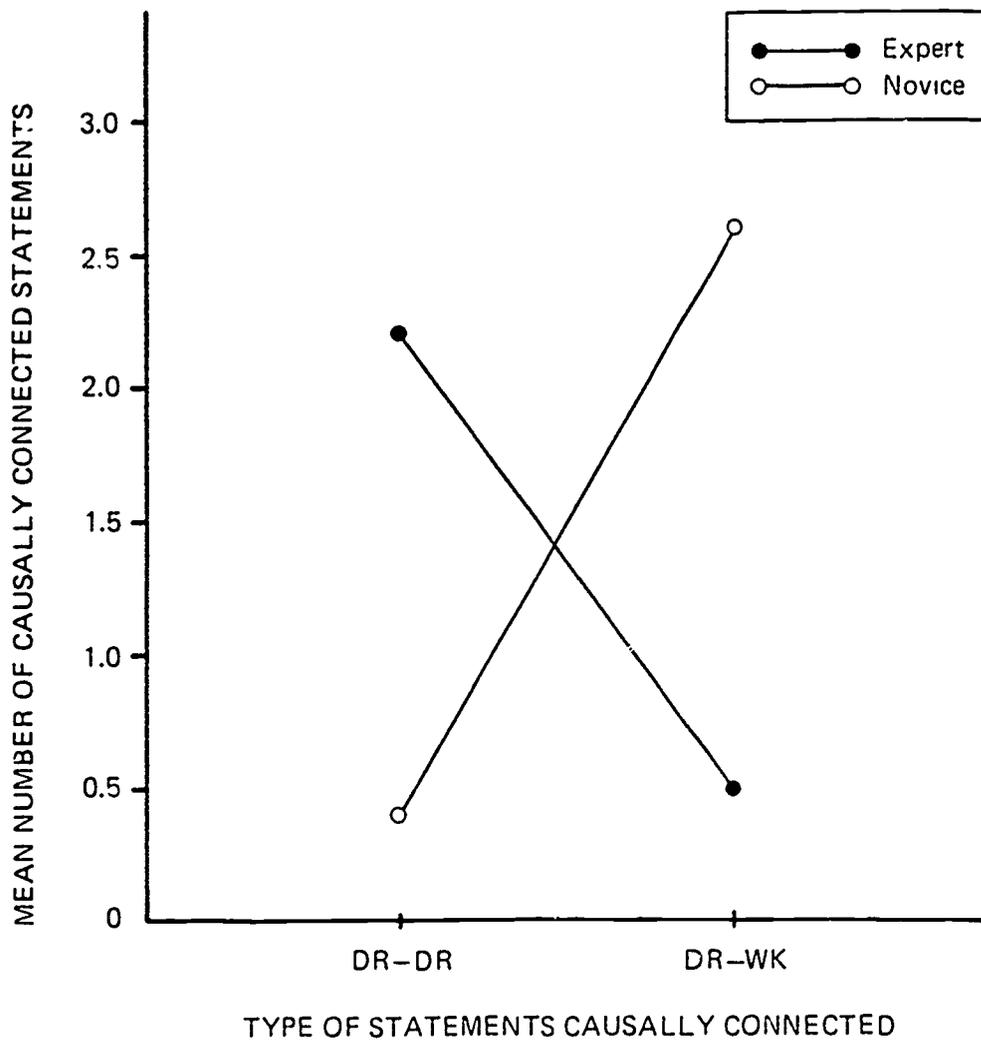


Figure 6

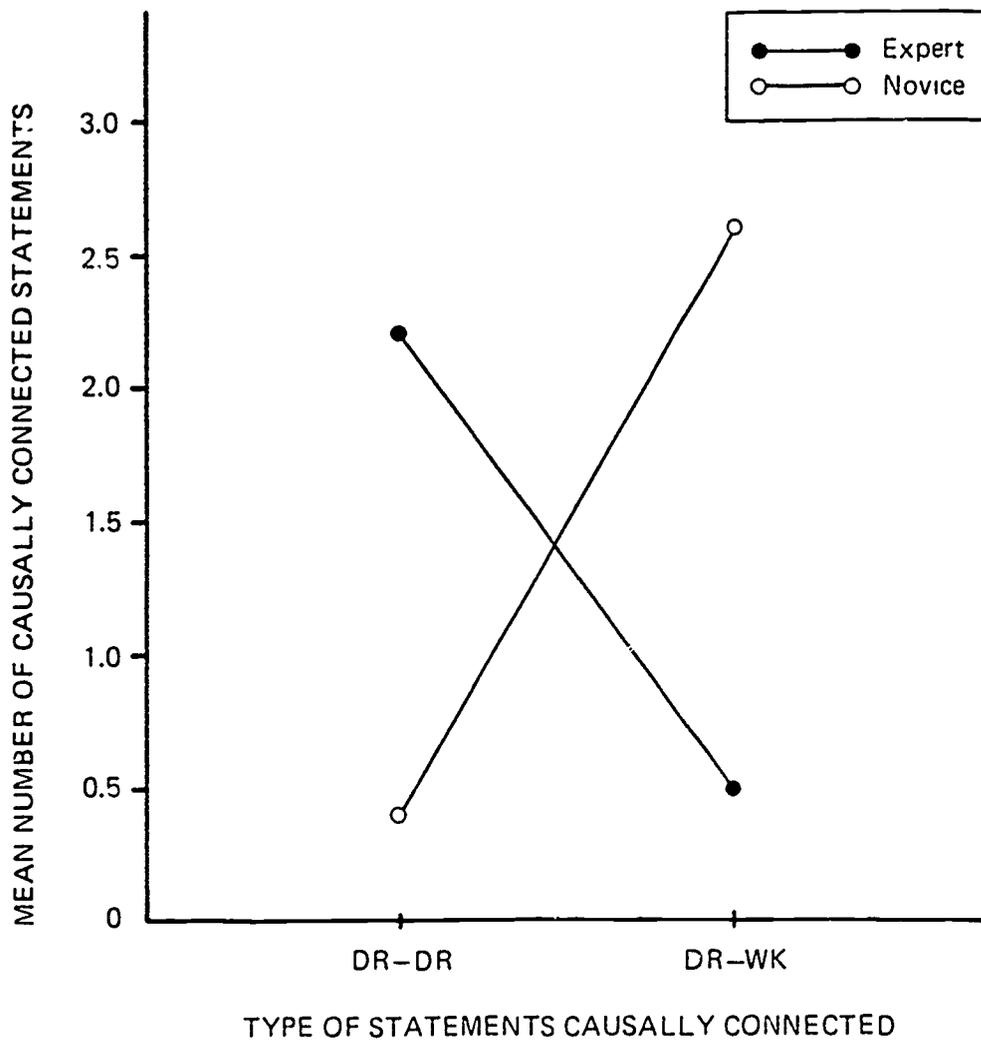


Figure 6