

Converging operations on a basic level in event taxonomies

MICHAEL W. MORRIS

University of Michigan, Ann Arbor, Michigan

and

GREGORY L. MURPHY

Brown University, Providence, Rhode Island

Research on object concepts has identified one level of abstraction as "basic" in cognition and communication. We investigated whether concepts for routine social events have a basic level by replicating the converging operations used to investigate object concepts. In Experiment 1, subjects were presented with event names from a taxonomy and were asked to list the actions comprising the event. Many more actions were listed at the middle than at the highest taxonomic level, without a further increase at the most specific level, paralleling the pattern of superordinate-, basic-, and subordinate-level object concepts. From these action lists, brief stories were composed for each event. In Experiment 2, subjects made pairwise similarity judgments on the stories. The mean similarity of events increased with specificity, as expected. But differentiation of categories (within-category similarity compared to between-category similarity) was highest for superordinates, contrary to results with object categories. In Experiment 3, subjects were fastest in recognizing actions as belonging to events named at the basic level. In Experiment 4, subjects predominantly chose basic-level terms to name stories. We conclude that event taxonomies do show basic-level structure, albeit a less sharply defined and less stable structure than in object taxonomies. The benefits and hazards of extending models of object concepts to other entities, such as social events, are discussed.

The question, "How shall a thing be called?" led Roger Brown (1958) to observe that out of the myriad words that can describe an object, one level of generality stands out as preferred. We call what we sit on a *chair* as opposed to a *desk chair* or a *piece of furniture*. The story tells of the *apple* that fell on Newton's head, not the *granny smith* or the piece of *fruit*. Brown concluded that we name objects at the level most useful in most contexts, using the name that generally distinguishes the referent from other objects.

Eleanor Rosch and her colleagues (Rosch, 1978; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) provided a more comprehensive answer with their proposal that objects are encoded in memory primarily at a medium level of abstraction, which they called the *basic* level of concepts, although more general *superordinate* and more specific *subordinate* concepts also play roles in cognition. Rosch et al. (1976) found that concepts at the basic level are optimally informative and show a preferred status in many cognitive processes, such as recognition and nam-

ing. In this paper, we investigate this principle of category structure beyond the domain of objects. Our experiments were designed to investigate whether concepts for events also have a basic level with similar processing advantages.

The basic level is a compromise. On the one hand, concepts should be specific and informative, so that knowing what category a thing belongs to will reveal many attributes. On the other hand, concepts should be general and inclusive, in order to minimize the total number of different categories a person needs. Basic-level concepts resolve this tradeoff by optimally combining distinctiveness and informativeness (Murphy & Brownell, 1985; Rosch et al., 1976). For example, consider the basic concept *table*. Tables are distinct from other types of furniture (they are not very similar to chairs, beds, lamps, bureaus, etc.), and knowing that something is a table gives important information about it (its shape, function, construction, etc.). The subordinate *coffee table* is much less distinctive, since coffee tables are very similar to other kinds of tables. The superordinate *furniture* is much less informative than the basic concept; knowing that something is a piece of furniture does not specify its shape, parts, function, or materials to any degree. Thus, the only level that is both distinctive and informative is the basic level.

Rosch et al. (1976) had two goals in their studies of the basic level. The first was to discover whether there is one level of abstraction that maximizes the combina-

This research was supported by NIMH Grant MH-41704 and NSF Grant BNS 83-15145 and was completed while the first author was at Brown University. We thank Jean Liittschwager for her invaluable help in running the experiments, and Ed Smith and the reviewers for their helpful comments on the manuscript. Correspondence may be addressed to Michael W. Morris, Department of Social Psychology, University of Michigan, 340 Thompson St., Ann Arbor, MI 48106.

tion of inclusiveness and informativeness. Their second goal was to demonstrate that basic concepts have a processing advantage (i.e., are used faster or more accurately) over other concepts in a variety of tasks.

Toward their first goal, Rosch et al. (1976) performed a number of feature-listing experiments. In their Experiment 1, subjects listed the attributes that most objects in a category have in common. In their Experiment 2, they collected the actions associated with objects in a category. The same pattern held in both experiments: subjects listed few features for abstract taxonomic terms, many features for the medium-level terms, and few additional features for the most specific terms. Tversky and Hemenway (1984) found qualitative as well as quantitative differences: distinctive features listed at the basic level are parts, whereas superordinate features are mainly functional, and subordinate features are mainly perceptual. In succeeding experiments, Rosch et al. fulfilled their second goal, finding that basic concepts were the most useful in a variety of perceptual and linguistic tasks (see below).

Since Rosch et al. (1976) outlined the basic-level theory and found evidence for it in the domain of objects, others have extended it to domains farther and farther removed from natural objects, including artificial stimuli (Mervis & Crisafi, 1982; Murphy & Smith, 1982), computer programming (e.g., sorting algorithms; Adelson, 1983), environmental scenes (e.g., beaches; Tversky & Hemenway, 1983), psychological situations (e.g., work; Cantor, Mischel, & Schwartz, 1982), and personality concepts (e.g., comic jokers; Cantor & Mischel, 1979).

Basic-Level Concepts for Events

Events can be named at various levels of abstraction, but as with objects, one level is generally preferred. For example, an acquaintance describing her plans for the day might report events like "taking a shower," "having breakfast," and "playing tennis." If she reported "performing hygienic activities," "eating a meal," and "playing a sport," she would sound evasive. Alternatively, if she reported "taking a cold shower," "having Wheaties," and "playing Canadian doubles," one might wonder what she was trying to emphasize with this extra detail (Cruse, 1977). Such intuitions suggest that there may be a basic level of event concepts.

Rosch (1978) was the first to suggest that the basic-level theory might be extended to concepts for events. She conducted a pilot study in which students were asked to list the events of their day. They also listed events for smaller and larger time units than days. The subjects reported events at a consistent level of abstraction. They described activities in terms of highly predictable sequences of actions, such as "feeding the cat," and not in more general terms, such as "all the morning chores," or more specific terms, such as "opening the cat food." Although suggestive, these results are inconclusive without the set of converging operations necessary to define a basic level (see Rosch et al., 1976). In particular, it is necessary to establish a consistent taxonomy and then compare event concepts from its different levels of abstraction.

The event concepts that Rosch tentatively called basic seem similar to the *scripts* for routine events, such as going to a restaurant, proposed by Schank and Abelson (1977). The research of Schank and his colleagues has primarily concerned the internal structure of scripts rather than the organization of event knowledge, and most of the psychological literature has followed this lead (Abbott, Black, & Smith, 1985; Bower, Black, & Turner, 1979). That is, rather than studying how scripts are related to each other and to more general event representations (i.e., studying taxonomies), they have investigated how the parts of individual scripts are related and used (see Barsalou & Sewell, 1985, for a discussion). Thus, although our project is related to research on script memory, it focuses on issues these researchers have not addressed.

One previous researcher, Rifkin (1985), systematically compared the levels of event taxonomies. Rifkin elicited features of event categories and found an apparent basic level. For compelling evidence, however, a series of converging operations, as in Rosch et al.'s (1976) original work, is required. To provide this evidence, we replicated experiments on object concepts using an extended version of Rifkin's taxonomies.

In Rosch et al.'s (1976) Experiment 2, subjects listed the motor movements associated with an object; subjects in our first experiment listed the actions that make up events. In this way, the behavioral informativeness of event terms can be measured at all three levels of the taxonomies. This experiment will also reveal the configuration of parts in event concepts, since actions are parts of events. Tversky and Hemenway (1984) showed that parts are the distinctive features of basic-level object concepts, and in Experiment 1 we investigated whether the same holds for events. In Experiment 2, we compared between-category and within-category similarity at each level of event concept (Mervis & Crisafi, 1982) as an additional measure of conceptual structure. Once the informativeness of event concepts is understood, we can then proceed to test for processing advantages. In Experiments 3 and 4, we investigated naming and categorization processes.

EXPERIMENT 1

Much of the information associated with objects or events is behavioral information. In Rosch et al.'s (1976) Experiment 2—behavioral measurement of category structure—subjects listed the motor movements associated with objects. In the current experiment, subjects listed the actions that make up events. We measured the behavioral informativeness of the three levels of event terms in Rifkin's (1985) taxonomies.

Method

Subjects. Thirty-five Brown University undergraduates were used as subjects. Most volunteered their time; a few were paid for their participation.

Materials. The stimuli were 45 taxonomically related event terms generated by Rifkin's (1985) subjects. There were nine general (i.e.,

superordinate) terms, with two terms at a medium level of detail for each general term and one specific term for each of the medium ones (see Rifkin, 1985, for details on how the categories were derived). This category structure has one drawback, namely that each middle-level event category (e.g., *wash hands*) has only one subcategory (e.g., *wash hands quickly*). Since some measures of category structure require two or more categories at each level, we used Rifkin's raw data (published in his Appendix B) to select a second specific event for every middle-level category. Out of the five or so candidates for each subordinate, we chose the category that was most distinctive from the one that Rifkin chose. The full set is shown in Table 1.

The stimuli were divided into seven sets of nine items each: one set of general terms, two sets of medium-level terms, and four sets of specific terms. Each set contained one item from each of the nine taxonomies, and 5 subjects responded to each set. Each subject received a booklet with instructions on the cover and one concept term at the top of each following page. The test booklets were assembled with the nine pages of items in a different random order for each subject.

Procedure. The subjects received test booklets and read the instructions. They were then allowed 2 min to respond to each event term. The instructions called event terms "activities," and the subjects were asked to list the parts of each activity. The parts of an activity were described as the "actions" common to all examples of that activity. The subjects were instructed not to write prose descriptions, but rather lists of short phrases describing actions,

and to avoid responding in terms of personal associations. The instructions contained two examples of activities and their parts. Depending on the level of events that a subject received, he or she saw the examples "going out to eat" and "vacation" (for those who received general terms), "restaurant" and "camping vacation" (middle-level terms), or "fast food restaurant" and "family camping vacation" (low-level terms). The example parts listed were the same across levels.

Judge amending. A judge (one of the authors) tallied all actions listed for each concept by at least 2 out of the 5 subjects who saw it. This criterion was similar to that used by Rosch et al. (1976). Making these tallies required equating similar actions listed by different subjects (e.g., "pay" and "pay for the meal"). To check the reliability of this process, a second judge (the other author) tallied the actions for five of the nine event taxonomies (the two highest, the two lowest, and the median taxonomy in number of actions tallied). There was high agreement ($r = .87$) between the action tallies of the two judges. Next, a new group of judges reviewed the action lists to amend our lists in the manner of Rosch et al. and Tversky and Hemenway's (1983) "judge-amended tallies." These previous researchers amended the lists to minimize artifactual results stemming from the pragmatics of the task. One possible artifact stems from subjects' following Grice's (1975) conversational maxim of relevance: Some actions listed for superordinate and basic-level events may be parts of the subordinate event as well, but they might have seemed too obvious for the subjects to mention. For example, "get a weapon" was listed for "Assault" but not for "Assault with

Table 1
Rifkin's (1985) Event Taxonomies

Superordinate	Basic	Subordinate
Meal	Breakfast	Quick Breakfast Going Out for Pancakes
	Dinner	Family Dinner Banquet Dinner
Entertainment	Movies	Horror Movies Western Movies
	Theatre	Comedy at the Theatre Drama at the Theatre
School Activities	Classes	English Class Math Class
	Tests	Essay Tests Multiple-Choice Exams
Sports	Football	Touch Football Tackle Football
	Hockey	Professional Hockey Olympic Hockey
Hygiene	Wash Hands	Wash Hands Quickly Rinsing Hands
	Shower	Lockerroom Shower Quick Shower
Shopping	Food Shopping	Shopping at a Butcher Shop Shopping for Fruit
	Furniture Shopping	Shopping for Bedroom Furniture Shopping for a Dining Room Set
Crime	Murder	Homocide Patricide
	Assault	Assault with a Deadly Weapon Mugging
Transportation	Train Travel	Subway Travel Amtrak to New York
	Bus Travel	Travel by City Bus Travel by School Bus
Housework	Vacuuming	Vacuuming the Carpet Vacuuming the Drapes
	Sweeping	Sweeping the Floor Sweeping a Rug

a Deadly Weapon," probably because it seemed already implied by the event name. Thus, the unamended results might have a bias that inflates the size of the basic-level advantage. Conversely, the items listed at the superordinate level could incorrectly include event instances instead of event parts, as Rifkin (1985) found. This would cause a bias that would deflate the basic-level advantage. Judge amending removed such biases.

Three naive graduate students and an author acted as judges. They were given 36 rating forms, each with one of the subordinate terms and its basic and superordinate category listed across the top of the page. All the actions given in the feature-listing task for the subordinate-level name and its dominant superordinate and basic terms were randomly listed down the left side of the page. The judges were instructed to decide whether each action was part of each activity. The instructions included Tversky and Hemenway's (1984) standard for the goodness of object parts, extending it to event parts: For an action to be a true part of an event, the judges were told, it must be perceptually discriminable, functionally distinct, and more or less obligatory to the activity. The judges also indicated when two of the listed actions were redundant; that is, when they referred to the same part of the event. The judges took the test booklets home with them and were allowed 2 days to complete them. Ratings were then compiled, with agreement of 3 out of the 4 judges as the criterion for including a part. The judges convened to resolve ratings of the items on which opinion was evenly divided. Items eliminated were either instances rather than parts (as *swimming* is to *sports*) or redundant actions. The final score for each category was its mean across all of the score sheets it appeared on.

Results and Discussion

On average, the subjects listed 6.1, 8.0, and 8.4 parts for the high-, middle-, and low-level categories, respectively. Although these results suggest a basic-level structure, such unamended means could be influenced by the biases mentioned above, so we will only analyze in detail the judge-amended means. We expected that the basic level would be the most general level in event taxonomies at which many parts are shared, and Table 2 reveals that this prediction was fulfilled. The three taxonomic levels received different numbers of features [$F(2,24) = 12.13$, $p < .001$]. For comparison with Rifkin's (1985) data, we first analyzed only the categories he studied. The middle-level event terms elicited significantly more parts than did the superordinates [$t(8) = 8.33$, $p < .001$]. The number of parts listed for middle-level event terms and for subordinate-level event names were nearly identical [$t(8) = 0.13$]. When the data for all 36 subordinates are included, the pattern is identical: the basic-level lists were significantly longer than the superordinate-level lists [$t(8) = 4.17$, $p < .001$], and the subordinate-level lists were not significantly longer than the basic-level lists [$t(8) = 0.19$]. In other words, as event names increase in generality from the subordinate level to the basic level, there is no loss in part information. But a further increase, to the more abstract superordinate terms, does create such a loss. Thus, these results agree with those of Rifkin in establishing a basic level of events.

One might be puzzled about the cases in Table 2 in which there are actually fewer actions listed for subordinate- than for basic-level events. In these cases,

Table 2
Number of Judge-Amended Actions Listed in Experiment 1

Concept	Level of Abstraction		
	Superordinate	Basic	Subordinate
Meal	10.5	14.2	15.0
Entertainment	2.0	16.0	16.0
School Activities	3.5	12.8	14.5
Sports	4.8	11.5	12.0
Hygiene	0.5	10.2	9.2
Shopping	11.8	16.2	16.5
Crime	1.8	5.0	4.5
Transportation	4.0	11.8	12.0
Housework	2.5	9.8	10.8
<i>M</i>	4.6	11.9	12.3

the judges decided that an action was generally true of the basic-level event but not of the subordinate. Rosch et al. (1976) used a forced inheritance rule in their taxonomies, such that any feature true of one category was assumed to be true of all lower level categories. Although feature inheritance is a property of all true taxonomies, when features and categories are not very typical this property does not hold in people's category representations. For example, although the subjects agreed that washing one's hands normally includes using soap, they also judged that washing one's hands quickly (a subordinate) does not require soap. Lack of feature inheritance is found in object categories as well (Hampton, 1987). We decided to accept the judges' decisions without forcing them into a taxonomic pattern of inheritance.

The number of features in both Rifkin's (1985) attribute-listing task and our action-listing task increases dramatically with increasing specificity to the basic level but not thereafter. One might hesitate to accept these experiments as independent, converging operations, since they are very similar. This objection can be answered by Rifkin's report that less than 14% of his subjects' features were classifiable as actions. So the results from Experiment 1 converge with Rifkin's to reveal basic-level structure in event concepts.

EXPERIMENT 2

Feature listing is a traditional method for describing the information in concepts. However, the sheer number of features listed for a concept cannot be taken as an exact measure of the information carried by that concept. One obvious problem with this operationalization is that there are known qualitative differences between the features listed at various levels of abstraction: functional features at the superordinate level, parts at the basic level, and perceptual features at the subordinate level (Tversky & Hemenway, 1984). Whether five functional features carry less information than six perceptual features is difficult to say. An alternative measure of information in a concept is the similarity of the category's members. Similarity

judgments may take into account the qualitative differences in information that are hidden in feature lists.

Experiment 2A Judging Event Stories

Mervis and Crisafi (1982) examined natural and artificial object concepts by comparing the similarity of category members at various levels of abstraction. Their subjects rated the similarity of pairs of objects, with the object pairs differing in taxonomic proximity. The pairs were from the same subordinate concept, from the same basic concept (but different subordinates), from the same superordinate (but different basic concepts), or from completely unrelated superordinates. Similarity increased from the general to more specific categories, but the most dramatic increase in similarity occurred between the superordinate and the basic levels. That is, two objects in the same basic category were much more similar than two objects merely in the same superordinate category. Being in the same subordinate category added little similarity over being in the same basic category, leading Mervis and Crisafi to the conclusion that basic concepts are the most differentiated. That is, only basic categories have both high within-category similarity and low between-category similarity. Superordinates lack the first and subordinates lack the second.

In Experiment 2A, we replicated Mervis and Crisafi's (1982) method, with one exception. Mervis and Crisafi used pictures of objects, but because one cannot take a snapshot of an event, we composed stories about our events based on subjects' action lists from Experiment 1. These stories were the stimuli for the similarity ratings. If there is a basic level of events, then the differentiation scores derived from this procedure should converge with the data of Experiment 1 to reveal the same level as being basic.

Method

Stimuli. Two brief stories were composed corresponding to each of the 36 subordinate-level event terms listed by the subjects in Experiment 1. The 72 stories were simple accounts of a character doing the event. Each action in the judge-amended lists (from Experiment 1) for a subordinate term became a clause in the story. The subordinate event term itself occurred in the first sentence of the story, since the subjects had listed actions in that context. For example, a story constructed for the subordinate *horror movies* follows:

Jean wanted to see a horror movie. She searched through the listings in the paper for a good film, then picked a show time and called some friends to go with her. Jean went to the theater and bought a ticket. She bought some food, and then gave her ticket to the usher as she entered the movie room. She picked a seat and sat down. Jean watched the preview and the show, screaming during the scary parts. When it was over, she left.

The two distinct stories for each subordinate were derived by varying the names of characters and places and by adding a unique action to each. These additional actions were drawn from the raw data in Experiment 1. They were actions listed for which there was insufficient consensus to be included in the final action list. This addition was necessary to prevent the subordinate stories from being completely identical.

Ninety story pairs were assembled: 36 same-subordinate pairs, which were simply the two stories derived from each subordinate; 18 same-basic pairs, which were stories from the same basic category but different subordinates; 18 same-superordinate pairs, which were stories from the same superordinate but different basic categories; and 18 unrelated pairs, which were stories from different superordinates. The particular pairings were determined randomly whenever possible. Each subordinate in the taxonomy appeared once in each condition (same-subordinate, same-basic, etc.). Each pair of paragraph-long stories was mounted on a large index card, and the cards were randomly numbered 1-90.

Subjects. The subjects were 16 Brown University undergraduates, who were paid for their participation.

Procedure. Each subject received half of the 90 story pairs, in one of these orders: 1-45, 45-1, 46-90, or 90-46. The subjects judged the similarity of the two stories using a 1 (not at all similar) to 9 (extremely similar) scale. The instructions emphasized that the subjects should read the stories thoroughly. The scale and instructions were the same as those used by Mervis and Crisafi (1982), allowing us to compare our results to theirs.

Results

The mean similarity of events of all four kinds was compared to Mervis and Crisafi's (1982) results for objects, and is presented in Table 3. At a gross level, the patterns are the same: similarity increased with increasing taxonomic proximity [$F(3,45) = 237.82, p < .001$]. This result is not very significant, however, since it shows only that the more specific categories were more similar, which was never in doubt. At a finer level, the patterns diverge. Superordinate event concepts seem coherent—their members are similar relative to the unrelated concepts. This difference became clearer when differentiation scores were calculated following Mervis and Crisafi's method. In that procedure, the similarity of objects within a category has the similarity of objects between categories at that level subtracted from it, providing a measure of within- versus between-category similarity at each category level. For example, the basic-level differentiation score is the mean basic similarity minus the mean superordinate similarity. Instead of finding the basic level to be the most differentiated, as Mervis and Crisafi found with objects (see Table 3), our results show the greatest differentiation at the superordinate level. The scores were 2.97, 1.56, and 1.71 for superordinate, basic, and subordinate categories, respectively. That is, superordinates showed the greatest difference between within-category

Table 3
Mean Rated Similarity of Event Pairs versus Object Pairs

Stimuli	Relationship			
	Same Subordinate	Same Basic	Same Superordinate	Unrelated
Event Stories	8.12	6.41	4.85	1.88
Event Names	8.01	6.58	5.15	1.97
Pictures of				
Natural Objects	7.70	6.91	3.27	1.18
Pictures of				
Artificial Objects	7.73	6.91	2.85	1.38

Note—The event data are from Experiment 2, and the object data presented for comparison are from Mervis and Crisafi (1982). Rating was done on a 1-9 scale.

similarity and between-category similarity. Almost as surprising is the finding that basic and subordinate categories have about the same differentiation scores.

Experiment 2B Judging Event Names

One objection to our methodology might be that we are measuring the superficial similarity of our stories, not the underlying similarity of events. That is, perhaps something in the way the stories were written, rather than the true similarities of the events, caused the pattern of results. To answer this criticism, we replicated Experiment 2A without using the stories, asking subjects to judge the categories only on the basis of their names.

Method

The subjects were 12 Brown University undergraduates, who were paid for their participation. The only difference between Experiments 2A and 2B was that in Experiment 2B, the subjects were given pairs of subordinate event names, such as *quick breakfast* and *family dinner*, and were asked to rate their similarity. Because they did not have to read entire stories, each subject rated all 90 pairs of names. This procedure created a problem for measuring the similarity within the lowest level categories, because there were no names at a still lower level that could be compared (i.e., no measure of within-category similarity). In order to replicate the previous experiment as closely as possible, we simply presented the subordinate name twice in this condition, asking the subjects on such trials to judge how similar all the events having that name were (e.g., how similar are all English classes?). Although this was a novel task, the subjects found it perfectly interpretable, and the results were quite regular. (In any case, the most important question being addressed was the relative differentiation of the basic and superordinate categories because of the unexpected result of the previous experiment. That comparison was not affected by the subordinate testing procedure.)

Results and Discussion

Table 3 shows the mean similarity ratings for event names. The different levels were again reliably different [$F(3,33) = 258.25, p < .001$]. The striking similarity of these results with the results from Experiment 2A demonstrates that the effect stems not from the surface properties of the stories but from the actual properties of the event concepts. The middle-level concepts did not maximize differentiation; instead, the superordinate level again was the best in this regard, with a score of 3.18, as compared to 1.43 for both the basic and the subordinate categories. Whether subjects read stories or simply used memory to judge similarity, the results were rather different from comparable measures of object concepts. In short, the results of Mervis and Crisafi (1982) were not replicated with events.

The similarity measures reveal that, despite the dearth of features listed for them in Experiment 1, superordinate concepts carry salient information. Note that this experiment did not disconfirm the predicted informativeness of middle-level categories. Rather, the results merely show that superordinates are considerably more informative than the feature listings indicate. We postpone interpretation of this finding until the General Discussion.

EXPERIMENT 3

Objects are often first seen or recognized as members of their basic-level category (Jolicoeur, Gluck, & Kosslyn, 1984; Murphy & Brownell, 1985; Smith, Balzano, & Walker, 1978). If events have a basic level, subjects should be able to verify category membership at that level more rapidly than at the superordinate or subordinate levels. This should be true even though superordinates have fewer features than basic concepts (which would seem to make verification of them easier). In Experiment 3, we tested this prediction for event concepts.

Unfortunately, it is impossible to use the same methodology for events as is used for object concepts. As in Experiment 2, we could not use photographs because, unlike objects, events cannot be depicted unambiguously in static pictures. Moreover, we could not use videotaped events or complete stories, because depicting a subordinate could require as many as 16 actions (see Table 2). In such stimuli, it would be difficult to prevent information about subordinate category membership from appearing later than information about more general categories. Therefore, the stimuli used were phrases describing parts of events rather than entire instances of events. The subjects judged whether an action (taken from the Experiment 1 lists) belonged to an event category—subordinate, basic, or superordinate. Although this paradigm deviates somewhat from the familiar object-categorization task, it should still reveal at which level events are recognized. In normal experience, events can be categorized as they unfold, on the basis of their parts rather than on the whole instance. Thus, we expected categorization reaction time (RT) to be fastest at the basic level.

Method

Materials. For each subordinate event, two actions were selected from the Experiment 1 action lists. The most distinctive actions were chosen for each event; we tried to find actions that are definitely part of that subordinate event but definitely not part of the three contrasting subordinates. For two of the nine taxonomies (*hygiene* and *sports*), nearly all of the actions of one subordinate also fit contrasting subordinates, so only the remaining seven taxonomies were used. However, it should be noted that some of the actions from one subordinate seemed plausible actions for the other subordinate and basic categories, even though the subjects in Experiment 1 had not listed them for those categories. (This problem arises because even though some actions are only typical or expected actions for one category, they are *possible* actions for similar categories. Thus, it becomes difficult to find indisputably false items.) This problem involves the construction of false trials, whereas the true trials were of primary interest. Each action was paired with its three "true" categories (one at each level) and with three categories that were most related while still being "false," resulting in 336 total items. Thus, for the action "scream during the scary parts," the true categories were *horror movie* (subordinate), *movie* (basic), and *entertainment* (superordinate), and the false categories were *western movie* (subordinate), *theatre* (basic), and *transportation* (superordinate). Because any of the superordinate categories (except the correct one) could be a false item, superordinates were randomly assigned to false actions such that each one occurred equally often.

The 336 items were randomly split into two lists of 168 items such that every category appeared in every condition once (thus,

all conditions appeared equally often within each list). Half of the subjects received one list and half received the other. The order of trials was randomly determined separately for each subject.

Subjects. The subjects were 16 Brown University undergraduates.

Procedure. The subjects were told that an event name would appear on the computer screen and that after they had read and understood it, they should press either one of two response buttons. The event name would disappear, and (200 msec later) an action would appear. The subjects were to decide whether that action was part of the named event. If so, they were to press the button marked "TRUE"; if not, they were to press the button marked "FALSE." The subjects used the dominant hand for the TRUE response and the nondominant hand for the FALSE response. After a 2-sec delay, the next trial began. Times under 200 msec or over 5 sec were deleted as outliers or buttonpress errors.

Results and Discussion

The true trials (shown in Table 4) led to the predicted results, but false trials were distorted by the problem (described above) of overlapping actions. For true trials, the subjects were fastest at the basic level, with the subordinate level close behind and superordinates quite a bit slower. The three levels were reliably different overall [$F(2,30) = 9.61, p < .001$ in the subject analysis, $F(2,54) = 4.29, p < .025$ in the item analysis]. The subordinate level was not significantly different from the basic level (both F s < 1), but the superordinate level was slower [$F(1,30) = 18.99, p < .001$, and $F(1,54) = 8.52, p < .01$]. Error rates increased slightly with category level [$F(2,30) = 7.40, p < .005$, and $F(2,54) = 4.49, p < .025$], although this may reflect the greater inclusiveness of higher level categories rather than a mistaken decision (i.e., they may not be errors). Overall, this is a pattern quite similar to results for object categories, in which superordinates are quite slow and subordinates are nearly as fast as basic categories (e.g., see Jolicoeur et al., 1984; Murphy & Brownell, 1985; Murphy & Smith, 1982).

For false trials, the most obvious result was that the subjects did not often agree that the item was false: only 43% and 51% of the subordinate and basic responses, respectively, were judged false. As noted above, this reflects the difficulty of finding truly distinctive features in the feature lists. Since superordinate items were in fact false 90% of the time, there were highly reliable error effects (though whether these judgments are really errors is very doubtful) [$F(2,30) = 199.81, p < .001$, and $F(2,52) = 32.18, p < .001$]. The false RTs show the corresponding pattern: long RTs for subordinate and basic

trials and short RTs for superordinates [1,610, 1,602, and 1,366 msec, respectively; $F(2,30) = 12.71, p < .001$, and $F(2,52) = 7.46, p < .005$]. (Because of missing data, one item had to be eliminated from the false RT analyses.) These slow RTs also undoubtedly reflect the fact that these actions were not clearly false for subordinate and basic concepts. In short, because of the problem in finding convincingly false items, the results from the false trials are probably not fully interpretable. Because so many of the false items were difficult (not clearly true or false), the subjects were probably more conservative in responding true (McCloskey & Glucksberg, 1979). There is no reason to believe that this could account for the observed basic-level advantage. In fact, since the false trials were easiest at the superordinate level, this should have aided true decisions at that level—because its true and false items would have been the most discriminable. Nonetheless, the superordinate true RTs were slowest of all.

The RT advantage of basics over subordinates in true trials was not statistically significant as it was in Rosch et al.'s (1976) results (Experiment 7). Yet, since other experiments with object concepts have shown that this effect occurs only under very particular conditions (Murphy & Brownell, 1985), this failure to replicate Rosch et al.'s result is unsurprising. The results of the present experiment, then, support the existence of a basic level of event concepts insofar as they are comparable with prior studies. The true trials followed the results of Murphy and Smith (1982), Jolicoeur et al. (1984), and Murphy and Brownell (1985). Because of the difficulty in creating false items, it is somewhat unfair to compare our false results to those of object-categorization experiments. However, this fact raises the rather puzzling question of *why* it was so difficult to devise false items for our stimuli (recall that this problem arose even after two hierarchies were eliminated because of the overlap problem). One possible reason is that in object-categorization experiments, a picture of the entire object is presented, whereas only one part of each event was presented in this experiment. Perhaps if object-identification tasks used single parts, such as *wings*, *seat*, or *handle*, similar results would obtain. We suggest other possibilities in the General Discussion.

EXPERIMENT 4

Expanding on Brown's (1958) observations, Rosch et al. (1976) found that basic-level names are almost universally preferred in object naming. They explained this striking effect by positing that basic concepts center on naturally occurring correlations of features in the perceived world, and argued that communication is easiest when terms are chosen that exploit these clusters of information. If the level of event terms identified by Experiment 1 is truly a basic level, these terms should be preferred as names for events.

Table 4
Latency and Accuracy of Categorization in Experiment 3

	Level of Abstraction		
	Superordinate	Basic	Subordinate
Mean Reaction Time	1,481	1,281	1,305
Percentage of "Yes" Judgments	79.7	84.8	89.5

Note—Mean reaction times are expressed in milliseconds.

Experiment 4A Naming Events

To study the terms most useful in communication, researchers have relied primarily on the free-naming task. In this procedure, subjects name pictured objects or scenes, and the terms they use are compared to names in the taxonomy. Although names at many levels of detail could be used, subjects consistently choose basic-level terms (e.g., see Berlin, Breedlove, & Raven, 1973). The basic-name preference is also found in other, more natural linguistic settings (Brown, 1958; Downing, 1980; Lucariello & Nelson, 1986; Wisniewski & Murphy, 1989). Once again, the picture-naming task used in earlier investigations was adapted in the present experiment to a story-naming task. The terms subjects chose to name the events in the stories was compared to the taxonomic event terms.

Basic-level names have been described as the terms of reference appropriate in neutral or average contexts (Cruse, 1977). However, names used in definite reference can change considerably with context: people often choose names at the level of detail sufficient to distinguish the referent from other things around—that is, from its contrast set (Olson, 1970). Somewhat surprisingly, Rosch et al. (1976, Experiment 10) found that basic category names are preferred across contexts—not just in neutral contexts. That is, even when all the objects in a stimulus set could be distinguished by superordinate terms, subjects used basic-level terms. And when subordinate names were necessary to distinguish each object, subjects still predominantly used basic category terms. To investigate whether shifting the context affects preference for basic-level names of events, we also compared three contrast sets.

Method

Stimuli. The 72 stories from Experiment 2 were used as stimuli. These stories were derived from the subordinate categories, which ensured that terms at all three levels could appropriately name the event described. Recall that the subordinate event term itself occurred in the first sentences of the story, since the subjects in Experiment 1 had listed actions in that context. Although this procedure may bias subjects toward using the subordinate name, it ensured that the event could be identified at that level.

The contrast set conditions were created by varying the similarity of the stories in the set. The superordinate contrast set contained no stories from the same superordinate categories, so the stories could be distinguished with superordinate terms. The basic contrast set contained stories from the same superordinate but not from the same basic categories, so that basic terms were sufficient to distinguish them. The subordinate contrast set contained stories from the same basic categories, so that subordinate terms were required to distinguish them. This manipulation was designed to discover whether subjects would use the most general name that would distinguish the stories (Olson, 1970) or instead use the basic level regardless of context, as Rosch et al. (1976) found. For the superordinate contrast condition, two sets of nine stories (one from each superordinate category) were created. For the basic contrast condition, one set of 18 stories (one from each basic category) was created. For the subordinate contrast condition, two sets were created—one with all 4 subordinate stories from each of the first 4 taxonomies (16 stories), the other with all 4 subordinates from

the remaining 5 taxonomies (20 stories). For all five sets, there were two versions: one in a random order and the other in the reverse order.

Subjects. The subjects were 40 Brown University undergraduates, who were paid to participate. Each set of stories was named by 8 subjects (4 received the forward version, and 4 received the reverse version), so that an equal number of subjects saw each story in each contrast set condition.

Procedure. The subjects were first instructed to read through the sets of stories quickly. This procedure ensured that they were aware of the contrast set before naming. They were then instructed to read through the set again, this time naming the event in each story with a short phrase. To discourage idiosyncratic, artistic titles, the instructions borrowed Tversky and Hemenway's (1983) injunction to give "the most simple, obvious, direct sort of name that ordinary people would give for each event."

Item tally. A judging procedure was used to classify responses as superordinate-, basic-, or subordinate-level names. Judges (an author and an independent judge) were given scoring sheets that listed the source subordinate term and its basic and superordinate terms for each story. The judges were instructed to "match each name on the answer sheet to one of the three on the scoring sheet ... choosing the name on the score sheet that resembles the subject's answer most closely in level of detail or generality." The two judges agreed on 90% of the names, and a third independent judge resolved the conflicts.

Results

This experiment was designed to see whether subjects would use basic-level names to label events, even when more specific or general names would be appropriate. Table 5 shows the number of names from each taxonomic level used by the subjects across contrast set conditions. Clearly, the basic-level names were favored, and superordinate names were hardly ever used. The differences between the three levels were reliable [$\chi^2(2) = 230.22$, $p < .001$]. A second question is whether contrast set affects the level of detail at which events are named. A test of independence found that the number of names at each level did depend on the contrast set [$\chi^2(4) = 31.83$, $p < .001$]. As Table 5 clearly shows, superordinate names were seldom used in any condition—even when they were sufficient to distinguish the stories. But subordinate names were used far more often when they were necessary to distinguish the stories than in other conditions. Thus, the basic-level advantage was not completely independent of contrast set, contrary to what Rosch et al. (1976) found for object names. Nonetheless, basic category names were used most often in all conditions, showing that the subjects were not simply choosing the names that would give an unambiguous reference (see Clark & Murphy, 1982).

Table 5
Number of Names Used at Each Category Level in Experiment 4

Contrast Set	Level of Abstraction		
	Superordinate	Basic	Subordinate
Superordinate	12	102	30
Basic	3	108	33
Subordinate	18	140	130
Total	33	350	193

Experiment 4B Rating Event Names

Before drawing any conclusions about the role of basic-level event concepts in communication, some alternative explanations must be ruled out. It may be that subjects avoid subordinate-level terms because they are unfamiliar. Another possible interpretation is that subjects avoid specific names because details of the stories deviate from their experiences. If a subject had attended horror movies only at drive-in theaters, he or she might consider "horror movies" an inappropriate name for the story about a horror movie at a walk-in theater. For that subject, the most specific name fitting the story would be the basic-level term "movies." To test both of these critical assumptions—that the subordinate terms are familiar to the subjects and appropriate to the stories—we conducted a final experiment in which subjects rated the appropriateness of all the taxonomic terms for the stories. We expected that the subordinate names would be rated at least as appropriate as the basic-level names, thus ruling out these alternative explanations for why subjects produced basic names in Experiment 4A. In fact, because the subordinate names describe the stories most precisely, they may be judged the most appropriate.

Method

After the subjects finished Experiment 4A, their answer sheets were collected, and they were instructed to read the stories again quickly. The subjects were then given rating sheets with the three possible names for each story in their set. Thus, corresponding to the story about going to a horror movie, there appeared:

Entertainment_____ Movies_____ Horror Movies_____

The instructions asked for ratings of the "appropriateness of each name for the event" on a scale of 1–10, where 1 means "not at all appropriate" and 10 means "most appropriate."

Results and Discussion

This experiment was designed to discover whether subjects consider subordinate-level terms appropriate names for the events in the stories, even though they chose basic-level terms more often in the free-naming task. In fact, the subjects rated subordinate terms as more appropriate names for the stories than basic-level terms. Table 6 shows mean appropriateness ratings for names at each level in each contrast set condition. Superordinate terms were rated much less appropriate than the other two levels, and the effect of category level was highly reliable [$F(2,16) = 157.65, p < .001$]. The subjects undoubtedly

felt that subordinate terms were the most appropriate names for the stories, but these were not the names spontaneously chosen in most conditions. Although there was a main effect of contrast set [$F(2,16) = 21.76, p < .001$], and the interaction between taxonomic level and contrast set was significant [$F(4,32) = 3.14, p < .05$], these effects do not mitigate the conclusions, since the subordinates were rated more appropriate than the basic terms in every contrast condition. These results justify the assumption that the subjects used basic-level names (in Experiment 4A) even when subordinate names would have been fully appropriate.

Experiment 4 showed that basic-level event names are greatly preferred in naming. Superordinates were virtually never used, and subordinates were used only in the subordinate contrast condition, where they were necessary to distinguish the stories. But even in this condition, the subjects used basic names about as often as subordinate names. Nonetheless, this shift with contrast set was one difference between event naming and object naming reported by Rosch et al. (1976). One possible explanation of this difference is that although objects tend to have familiar nouns that denote their basic-level name, many events require multiple-word descriptions (e.g., *going to the theater, seeing a play*). Often these descriptions are not standardized, in that any of a number of phrases could be used, with none having clear priority. Perhaps the contrast set has a stronger influence on event naming because of the lack of conventional, standard names. Overall, event terms at the basic level are clearly the preferred name in neutral settings, and their superiority is only somewhat abated in other contexts.

The rules of cooperative communication may explain why the structural characteristics of the basic level lead subjects to prefer basic category names (Clark & Murphy, 1982; Cruse, 1977). One commonly accepted conversational rule is for speakers to be informative but not to give unnecessary information (Grice, 1975). Labeling an event with a superordinate name (e.g., *entertainment*) provides insufficient information about the event. As Table 2 shows, most of the concrete information about an event is given by the basic level. Superordinate names may be used when referring to a number of events at once (i.e., no single event may be called "entertainment," but a collection of events might be; Wisniewski & Murphy, 1989). Subordinate names may only be used when the specific information associated with that level is particularly relevant (Clark & Murphy, 1982). Even when the subordinate was necessary to distinguish an event from its contrast set, it was not the name that most subjects chose, probably because the marginal gain in information from moving to the subordinate was so small.

Unlike the naming results, appropriateness ratings were a direct function of the specificity of the name, with subordinates receiving the highest rating and superordinates receiving the lowest ratings. This pattern is very similar to that found in object concepts by Smith and Osherson (1984; Smith, Osherson, Rips, & Keane, 1988). They

Table 6
Mean Appropriateness Ratings from Experiment 4

Contrast Set	Level of Abstraction		
	Superordinate	Basic	Subordinate
Superordinate	3.93	7.66	9.12
Basic	4.52	8.53	9.47
Subordinate	4.96	7.88	9.64
<i>M</i>	4.47	8.02	9.41

found that an object was judged most typical of the most specific category that it fit (e.g., a red apple would be judged more typical of *red apple* than of *apple*). Their theory explains this result as a function of the number of shared features in the object and concept representation. That is, a red apple shares more features with the concept *red apple* than it does with *apple* or *fruit*. A similar account would explain the present results as well: The more specific an event concept, the more closely it will match the story, and the higher the appropriateness rating will be. Most importantly, the processing advantages of basic-level concepts cannot be explained away as reflecting only appropriateness or typicality.

GENERAL DISCUSSION

In this study, we repeated procedures used to identify the basic level of object concepts to investigate whether concepts for routine social events have basic-level structure. The general pattern of results from our four experiments seems to replicate the pattern of informativeness and processing advantages marking the basic level. Against the general trend of convergence, however, were two notable divergences: the finding that superordinate event concepts were very differentiated in similarity ratings (Experiment 2) and the shift in names chosen toward the subordinate level in contrast sets consisting of highly similar events (Experiment 4). One purpose in studying concepts for entities other than material objects is to isolate the properties of basic-level structure from the properties of the object domain. We interpret the similarities in results for objects and events as evidence for a common organizational structure—the basic level—across domains, and we interpret the differences as indications of how unique properties of each domain interact with the common structure.

Extending Rosch's method and theory from the object domain to entities in other ontological domains allows us to extend some of her general conclusions. There are two that we would like to highlight before discussing the details of event concepts. First, it seems that the structure of the world of events (as perceived by human agents) affords optimal categorization at an intermediate level of abstraction. This is because attributes are not distributed orthogonally or randomly across events; many attributes covary, and categories at an intermediate level can center on clusters of co-occurring attributes. Second, social event concepts serve cognitive purposes, as do object concepts, such that greater differentiation leads to an advantage in processing. These general conclusions are supported by the pattern of specific results described in the next section.

Event Representations

To review the interaction of basic-level organization and the unique properties of events, we propose an account of how event information, from abstract to concrete, is represented in concepts. To begin our discussion, let us first describe superordinate event concepts (e.g., *enter-*

tainment). We suggest that, like object superordinates, superordinates carry information about the functions, purposes, and goals of activities (Tversky & Hemenway, 1984). Categorizing at this level highlights abstract commonalities between activities that have a similar role in our culture. The coherence of these concepts comes not from below, the sharing of perceivable features, but from above, their common role in "theories" of human behavior (Murphy & Medin, 1985). Hence, few features are given for superordinates in lists of attributes (Rifkin, 1985) or lists of actions (Experiment 1). Yet, because of their similar functional roles, these categories are coherent: events within the same superordinate are far more similar to each other than to events from different superordinates (Experiment 2), even though they do not share the same parts. For example, participants in all forms of entertainment share common goals and feelings, even if they differ in their actual actions. Since superordinate terms are not directly associated with specific actions, they are slower in the categorization task (Experiment 3). Finally, because superordinates do not provide concrete details about individual events, they are not used as names; more specific concepts are so used (Experiment 4).

Basic-level event concepts (e.g., *going to the movies*) carry the abstract functional information of superordinates but also a wealth of concrete, perceptual information. The critical features that make the basic level so primary may be *parts* (Tversky & Hemenway, 1984), which are functional and perceptual divisions of the event. When subjects list the parts of events (actions), the informativeness of basic-level concepts is dramatically revealed (Experiment 1). Other kinds of features such as objects, locations, times, and outcomes, also cluster at the basic level (Rifkin, 1985). This information makes basic concepts more coherent than superordinates (Experiment 2). Because basic-level concepts are strongly associated with specific actions, they are used more quickly in categorization than are concepts at other levels (Experiment 3). And because basic-level concepts are both distinctive and informative, they are generally chosen to name events (Experiment 4). The informativeness and processing advantages shown by concepts at the basic level suggest that memory for events may be primarily encoded at this level. In other words, these concepts are more available than more abstract or more concrete alternatives.

Individually, subordinate event concepts (e.g., *going to a horror movie*) are much like basic concepts with a bit more detail. Collectively, they differ in that they are more similar to each other, giving them little distinctiveness. That is, going to a horror movie is very much like going to a western. Furthermore, it seems likely that the distinct features of subordinates tend not to be parts of the event, but are perceptual and contextual details. For example, Experiment 1 found that, on average, subordinates added less than one new action to the actions that make up basic categories. Yet subordinates were considered by the subjects in Experiment 2 to have considerably more within-category similarity than basic categories,

so there must be other aspects (besides parts) of subordinates that make them similar. These aspects may include the manner of carrying out an action (as for *quick breakfast*), the objects involved (*assault with a deadly weapon*), or the location (*Amtrak to New York*), rather than entire parts. In categorization tasks, subordinates can be used no more quickly than basic concepts (Experiment 3), presumably because of their lack of distinctiveness. Subordinates are not preferred in naming generally, but in contexts in which their details are relevant for distinguishing events, subjects do use their names (Experiment 4).

Differences Between Object and Event Concepts

The description of event concepts that we have just proposed is in many respects quite similar to theories of levels of object concepts. Although there are some important caveats in this correspondence, it seems that some aspects of conceptual structure are universal enough to apply both to events and to objects—two very different ontological types. Nonetheless, there are some important differences between the two, and the rest of the discussion focuses on these differences.

One difference between object and event categories is the apparently different representations of superordinates. Experiment 2 found that event superordinates are highly informative, although in Experiment 1 few features were listed for them. The key to this paradox may lie in qualitative rather than quantitative differences between features. Object superordinates carry functional, not perceptual, information (Tversky & Hemenway, 1984). And Rifkin (1985, p. 542) noted that event superordinates contain a higher proportion of “abstract person-related features,” such as “competition” and “fun” for *sports*. These abstract features seem more central to event concepts than to object concepts, since events are structured by goals rather than by a perceptual shape. For example, what makes an activity entertainment is not some specified set of physical actions, but rather the fact that the participants are doing it in order to have fun. Similarly, Read, Jones, and Miller (in press) have argued that goals are central to personality concepts, and Hampson, John, and Goldberg (1986) have found that personality superordinates are relatively more informative than object superordinates.

We noted above that events often do not have ready-made names for them, as objects do. When someone asks you what you are doing, there is often no single name that is the conventional label for that activity. One might easily respond with a number of names that focus on different aspects of the activity, at different levels of abstraction and including more or fewer actions. This might explain why, although there is a basic-level naming advantage (Experiment 2), it is not as stable across contexts as in object hierarchies.

This difference between the codabilities of objects and events could be an artifact of the English language, but there is reason to believe that event categorization is intrinsically less stable than object categorization. Gentner (1982) analyzed the componentiality and boundedness of

actions versus objects in order to explain differences between nouns and verbs. Languages differ greatly in what aspects of an action are included in a single verb, whereas they largely agree on what counts as an individual noun referent. This occurs, Gentner argued, because actions have a number of components that are in fact separated in different actions. An action might include an actor, an object, a recipient, and an activity, done in a certain manner at a certain time in a certain location or context. However, at other times, the same activity might be done with a different actor, in a different manner, and so on. Different verbs (and different languages) encode different components of these actions, which suggests that there is no single, preferred way to conceptualize every action. The parts of objects, on the other hand, are often not physically separable from each other and do not occur alone. Equally important, the boundaries of objects are much more clearly defined than those of actions. It is often very difficult to delineate precisely when an event has begun and ended. Does a party begin with the preparations, the arrival of the first guest, or when all the guests have arrived? Parties, like many events, often have no well-marked beginning and ending. In contrast, there is little difficulty in discerning where a chair ends and the floor begins; the boundaries of objects are usually very clear.

Another reason that object categories have greater stability is that for some kinds of objects, physical laws impose a categorical structure. For example, *natural kinds*, such as chemical elements or biological species, are shaped by natural processes and seem to exist as categories prior to human activities (Putnam, 1970). In contrast, categories of social events are necessarily shaped by human goals. Event categories tend to be more heterogeneous and overlapping because human goals are plastic and multiplex. Vallacher and Wegner (1985, 1987) have analyzed how actions can be identified on many different teleological levels. For example, a single action can be categorized as *pressing buttons*, *making a phone call*, or *initiating a friendship*. Thus, the conceptualization of actions is more variable. This variability was seen in Experiment 3, where it was difficult to find actions that were true of one category but clearly false of related categories. However, this is not to say that there is no variability in object categorization, as Barsalou's (1983, 1987) work has demonstrated.

We have emphasized that there may be many different ways of classifying events. Cognitive psychologists often assume a taxonomic hierarchy for concepts because of its property of inferential inheritance (Collins & Quillian, 1969)—that is, properties of high-level categories are also properties of their subordinates. However, other students of events have focused on entirely different classification systems, such as partonomies (Abbott et al., 1985; Reiser, Black, & Abelson, 1985). For example, they might compare *going to a restaurant* with *ordering food* and *telling the waiter your selection*. And Vallacher and Wegner (1985, 1987) have discussed event concepts as organized by networks of goals. Script organizations, based on the

work of Schank and Abelson (1977), do not support inferential inheritance (e.g., if going to a restaurant makes you less hungry, it does not follow that the act of ordering makes you less hungry). Nonetheless, these organizations are crucial to the planning of activities and the understanding of the course of events in some activity. So an ordinary understanding of events probably includes a partonomic structure (as in script theory), a hierarchical structure (as studied in this article), and various goal-oriented relations. Now that the basic level in taxonomies has been identified, the question of whether it plays a special role in these other organizations can be addressed.

REFERENCES

- ABBOTT, V., BLACK, J. B., & SMITH, E. E. (1985). The representation of scripts in memory. *Journal of Memory & Language*, *24*, 179-199.
- ADELSON, B. (1983). *Constructs and phenomena common to semantically-rich domains* (Cognitive Science Tech. Rep. No. 14). New Haven, CT: Yale University.
- BARSALOU, L. W. (1983). Ad hoc categories. *Memory & Cognition*, *11*, 211-227.
- BARSALOU, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- BARSALOU, L. W., & SEWELL, D. R. (1985). Contrasting the representations of scripts and categories. *Journal of Memory & Language*, *24*, 646-665.
- BERLIN, B., BREEDLOVE, D. E., & RAVEN, P. H. (1973). General principles of classification and nomenclature in folk biology. *American Anthropologist*, *75*, 214-242.
- BOWER, G. H., BLACK, J. B., & TURNER, T. J. (1979). Scripts in memory for text. *Cognitive Psychology*, *11*, 177-220.
- BROWN, R. (1958). How shall a thing be called? *Psychological Review*, *65*, 14-21.
- CANTOR, N., & MISCHEL, W. (1979). Prototypes in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 12, pp. 3-52). New York: Academic Press.
- CANTOR, N., MISCHEL, W., & SCHWARTZ, J. C. (1982). A prototype analysis of psychological situations. *Cognitive Psychology*, *14*, 45-77.
- CLARK, H. H., & MURPHY, G. L. (1982). Audience design in meaning and reference. In J. F. Le Ny & W. Kintsch (Eds.), *Language and comprehension* (pp. 287-299). Amsterdam: North-Holland.
- COLLINS, A., & QUILLIAN, M. R. (1969). Retrieval time and semantic memory. *Journal of Verbal Learning & Verbal Behavior*, *8*, 240-247.
- CRUSE, D. A. (1977). The pragmatics of lexical specificity. *Journal of Linguistics*, *13*, 153-164.
- DOWNING, P. (1980). Factors influencing lexical choice in narrative. In W. L. Chafe (Ed.), *The pear stories: Cognitive, cultural, and linguistic aspects of narrative production* (pp. 89-126). Norwood, NJ: Ablex.
- GENTNER, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. A. Kuczaj II (Ed.), *Language development: 2. Language, thought and culture* (pp. 301-304). Hillsdale, NJ: Erlbaum.
- GRICE, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: 3. Speech acts* (pp. 41-58). New York: Academic Press.
- HAMPSON, S. E., JOHN, O. P., & GOLDBERG, L. R. (1986). Category breadth and hierarchical structure in personality: Studies in asymmetries in judgments of trait implications. *Journal of Personality & Social Psychology*, *51*, 37-54.
- HAMPTON, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, *15*, 55-71.
- JOLICOEUR, P., GLUCK, M., & KOSSLYN, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, *16*, 243-275.
- LUCARIELLO, J., & NELSON, K. (1986). Context effects on lexical specificity in maternal and child discourse. *Journal of Child Language*, *13*, 507-522.
- MCCLOSKEY, M., & GLUCKSBERG, S. (1979). Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cognitive Psychology*, *11*, 1-37.
- MERVIS, C. B., & CRISAFI, M. A. (1982). Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development*, *53*, 258-266.
- MURPHY, G. L., & BROWNELL, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *11*, 70-84.
- MURPHY, G. L., & MEDIN, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289-316.
- MURPHY, G. L., & SMITH, E. E. (1982). Basic-level superiority in picture categorization. *Journal of Verbal Learning & Verbal Behavior*, *21*, 1-20.
- OLSON, D. R. (1970). Language and thought: Aspects of a cognitive theory of semantics. *Psychological Review*, *77*, 257-273.
- PUTNAM, H. (1970). Is semantics possible? In H. E. Kiefer & M. K. Munitz (Eds.), *Language, belief, and metaphysics* (pp. 50-63). Albany, NY: State University of New York Press.
- READ, S. J., JONES, D. K., & MILLER, L. C. (in press). Traits as goal-based categories: The importance of goals in the coherence of dispositional categories. *Journal of Personality & Social Psychology*.
- REISER, B. D., BLACK, J. B., & ABELSON, R. P. (1985). Knowledge structures in the organization and retrieval of autobiographical memories. *Cognitive Psychology*, *17*, 89-137.
- RIFKIN, A. (1985). Evidence for a basic level in event taxonomies. *Memory & Cognition*, *13*, 538-556.
- ROSCH, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.
- ROSCH, E., MERVIS, C. B., GRAY, W., JOHNSON, D., & BOYEBRAEM, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382-439.
- SCHANK, R. C., & ABELSON, R. (1977). *Scripts, goals, plans and understanding*. Hillsdale, NJ: Erlbaum.
- SMITH, E. E., BALZANO, G. J., & WALKER, J. (1978). Nominal, perceptual, and semantic codes in picture categorization. In J. W. Cotton & R. J. Klatzky (Eds.), *Semantic factors in cognition* (pp. 137-168). Hillsdale, NJ: Erlbaum.
- SMITH, E. E., & OSHERSON, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science*, *8*, 337-361.
- SMITH, E. E., OSHERSON, D. N., RIPS, L. J., & KEANE, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, *12*, 485-527.
- TVERSKY, B., & HEMENWAY, K. (1983). Categories of environmental scenes. *Cognitive Psychology*, *15*, 121-149.
- TVERSKY, B., & HEMENWAY, K. (1984). Objects, parts and categories. *Journal of Experimental Psychology: General*, *113*, 169-193.
- VALLACHER, R. R., & WEGNER, D. M. (1985). *A theory of action identification*. Hillsdale, NJ: Erlbaum.
- VALLACHER, R. R., & WEGNER, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, *94*, 3-15.
- WISNIEWSKI, E. J., & MURPHY, G. L. (1989). Superordinate and basic category names in discourse: A textual analysis. *Discourse Processes*, *12*, 245-261.