MEMORY, 1997, 5 (5), 569–599

Multiple Organisations of Events in Memory

Juliana S. Lancaster
Morris Brown College, Atlanta, USA

Lawrence W. Barsalou
Emory University, Atlanta, USA

Theories of memory organisation propose that activity knowledge organises autobiographical memory globally. According to these views, memories that share a participant, location, or time are only organised together if they also share an activity. If they do not, they are nested within their respective activity organisations locally rather than being organised together globally. Two experiments that assessed people's clustering of laboratory events consistently obtained findings that contradict this view. Both experiments found that people organise event memories globally in non-activity clusters, cross-classify events into multiple organisations, and pivot between activity and non-activity clusters. Consistent with studies of naturalistic events, these studies of laboratory events indicate that people cross-classify event memories simultaneously into multiple global organisations.

INTRODUCTION

Memory for events has received increasing attention from memory researchers, as much recent work illustrates (e.g. for reviews, see Conway, 1990b; Neisser, 1982; Neisser & Winograd, 1988; Nelson, 1986; Rubin, 1986). Although this work has addressed numerous important aspects of autobiographical memory, it has provided little information about its organisation. To the extent that organisation has received attention, one account has dominated. According to this view, knowledge of activity types organises event memories as a by-product of the comprehension process (Kolodner, 1978, 1980, 1983a,b, 1984; Reiser, 1983, 1986; Reiser, Black, & Abelson, 1985; Reiser, Black, & Kalamarides, 1987; Schank, 1982; Schank & Kolodner, 1979; also see Barsalou, 1995).

Requests for reprints should be sent to Lawrence W. Barsalou, Department of Psychology, Emory University, Atlanta, GA 30322, USA. Email: barsalou@emory.edu

This article reports a subset of the experiments from the first author's doctoral dissertation performed at Emory University (Lancaster, 1985). Grants IST-808984, IRI-8609187, and SBR-9421326 from the National Science Foundation to the second author supported this work. We are grateful to Daniel Sewell for his contributions to this project, and to Martin Conway and two anonymous reviewers for their helpful comments on an earlier draft of this paper.

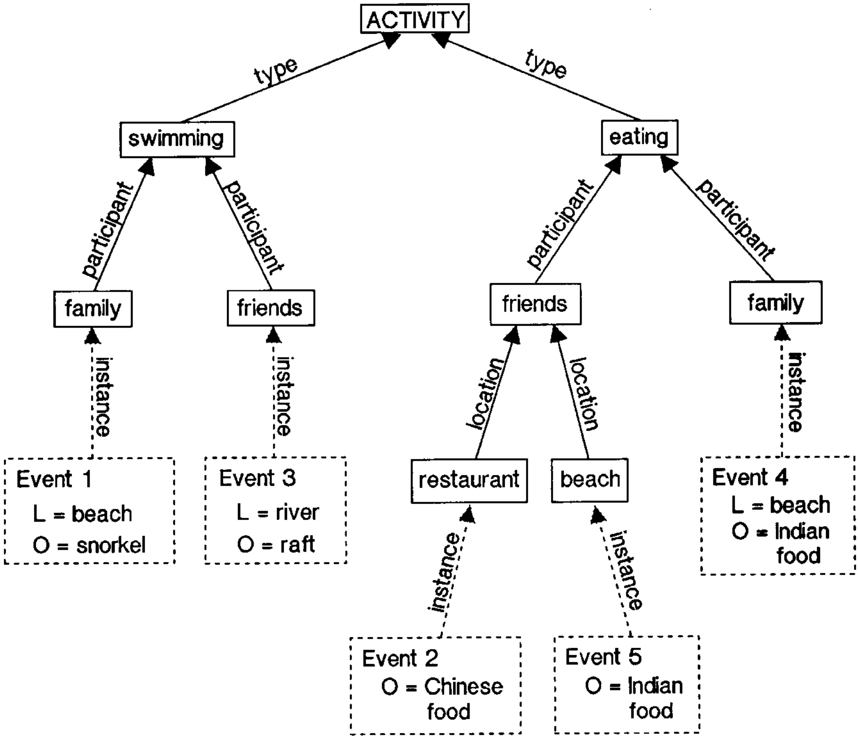

FIG 1. An example of event organisation according to the strong activity view. Only activities organise events globally, with other organisations, such as participants and locations, embedded locally. L is location, and O is object.

Because events involving the same activity are comprehended with the same abstract knowledge, memories of them become integrated together with it. For example, all events that involve eating at a restaurant are stored together with abstract knowledge about eating at restaurants; all events that involve going to movies are stored together with abstract knowledge about going to movies; and so forth. This abstract knowledge could take the form of scripts (Schank & Abelson, 1977), MOPs (Schank, 1982), E-MOPs (Kolodner, 1978, 1980, 1983a,b, 1984), or contexts (Reiser et al., 1985; Reiser et al., 1987).

Within one of these abstract forms of knowledge, event memories that share the same activity are organised by differences among them (e.g. participants, locations). For example, different memories about eating at restaurants might be further organised into participant subclusters, such as eating at restaurants with family, with friends, with co-workers, and so forth. In turn, these clusters could contain still more specific subclusters, such as clusters of events for eating Chinese food with friends, eating Indian food with friends, and so forth.

Figure 1 shows a set of events organised in this manner. As can be seen, activities provide the highest level of event organisation, with events being grouped most globally into those that involve swimming, those that involve eating, and so forth. Other global organisations of events do not occur. For example, events sharing the same participant (e.g. family) are not organised together but instead are distributed in two different activity organisations, with separate subclusters of events developing around participants in each activity organisation. For example, one subcluster of events that involves friends has formed under swimming, and another has formed under eating. With further events, still more specific subclusters could occur, such as eating with friends at restaurants.

Importantly, these subclusters are defined locally with respect to particular activities. For example, the two clusters of events that involve friends are local with respect to swimming and eating. There is no single cluster of events that involves friends independent of the activity. Consequently, this view predicts that a cluster of events sharing a participant should always share an activity as well. In retrieving events about friends, for example, people might only retrieve events that involve eating and fail to retrieve events that involve any other activity, such as swimming.

Retrieving an event memory from such an organisation first requires identifying the activity. If the retrieval cue specifies an activity (e.g. ‘‘eating’’ in the cue ‘‘eating Chinese food with friends’’), search becomes constrained to the cluster of events that all involve eating. If a retrieval cue fails to specify an activity (e.g. recall an event at the beach with family), then inference strategies identify candidate activities to search (Kolodner, 1984; Reiser et al., 1985). For example, swimming would be a reasonable activity to infer for an event at the beach. On inferring an activity, known aspects of the event (e.g. participants, location, objects) guide search through the memories of the activity, specifying the path to take through the activity organisation. On accessing the cluster of events that share swimming, for example, search might seek a subcluster for beach or family. When such details are missing and thereby prevent further search towards an event memory, various cue elaboration strategies provide plausible inferences about what the details might be (Kolodner, 1983b; Reiser et al., 1987). Once an event memory is found, search ends. If the initially selected activity does not produce a successful retrieval, these search and inference strategies proceed iteratively through other activity organisations. For example, if the sought-after event at the beach is not found under swimming, it might be searched for under fishing.

Importantly, these accounts of activity organisation predict that people cannot begin to search for an event memory without first having an activity cue that specifies the activity organisation to search. A particular location or participant cannot provide direct access to a sought-after event. Instead, such cues provide a means of inferring a probable activity organisation to search, and search can only begin once an activity has been selected.

In contrast, other organisational frameworks make different predictions. One alternative scheme assumes that event memories are cross-classified simultaneously in a variety of global organisations, including organisations for location, participant, time, and objects—as well as for activity (e.g. Barsalou, 1988). Figure 2 shows the same set of events as in Fig. 1 organised in this alternative manner. As can be seen, any event can be accessed directly by a number of retrieval cues. For example, Event 1 can be accessed directly from *family*, *beach*, or *snorkel*—it is not necessary to know or infer an activity initially. Furthermore, global clusters of events develop for other attributes besides activity. For example, all memories that involve friends in Fig. 2 are organised into a single global cluster. (Recall that these memories were distributed in two unrelated subclusters in Fig. 1.) Similarly, global clusters exist in Fig. 2 for family, beach, and Indian food. In general, this cross-classification approach to event organisation assumes that multiple organisations provide direct access to events and establish a wide variety of global event clusters.[1]

Recent findings suggest that people use other conceptual structures besides activities to organise events globally (for reviews see Conway, 1990b; Conway & Rubin, 1993). For example, people sometimes organise events globally with goal-derived categories (Conway, 1990a) or emotion concepts (Conway, 1990). On other occasions, people organise events with various temporal structures and schemata (Anderson & Conway, 1993; Eldridge, Barnard, & Bekerian, 1994). Several investigators have found that extended temporal events are particularly important in global event organisation (Barsalou, 1988; Brown, Shevell, & Rips, 1986; Conway & Bekerian, 1987). For example, people organise events according to extended stays in locations (e.g. when I lived in California), extended periods of schooling (e.g. when I was in college), and extended personal relationships (e.g. my second marriage). Extended events differ from activity types in several ways (Barsalou, 1988). For example, event memories are organised chronologically as parts of a single extended event rather than being organised conceptually as multiple, temporally unrelated instances of an activity type. Together, these findings suggest that non-activity organisations can organise event memories globally.

Finally, an intermediate view is that people can cross-classify events in multiple organisations globally but that activity organisations provide the dominant form of organisation. Following the intuitions behind the activity view of Kolodner, Reiser, and Schank, activities provide a prominent

[1]Figure 2 omits much important structure about the organisation of conceptual domains. For example, Fig. 2 groups *raft* together with *Indian* and *Chinese food*, failing to acknowledge that the two types of *food* form a conceptual subcluster. As Barsalou (1988) illustrates, however, the proper way to represent these conceptual domains is hierarchically. For example, objects might be divided into animate and inanimate objects, which might be further divided into still more specific categories, and so forth. Depending on how entities are categorised in an event, an event memory may point to objects at different levels in a hierarchy.

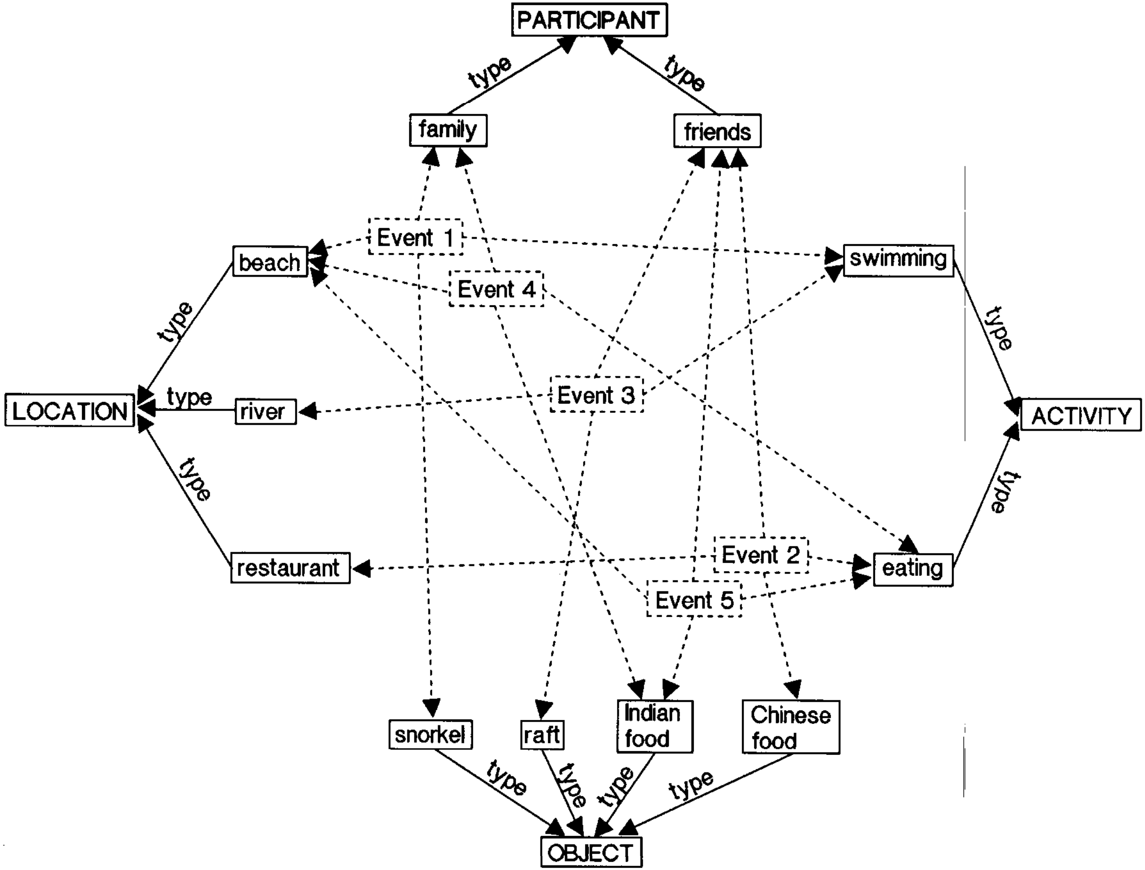

FIG. 2. An example of event organisation according to the cross-classification view. The same events as in Fig. 1 are organised globally in multiple organisations. Events 2 and 5 form a high-similarity cluster sharing an activity, whereas Events 1 and 4 form a high-similarity cluster sharing no activity.

organisation of events. Because activities are so central to comprehending events, they become highly integrated with event memories. However, because events are also processed with respect to their locations, participants, times, and objects, they also become directly integrated with these other conceptual domains (Barsalou, 1988, 1995). Nevertheless, if activities typically receive more processing than these other domains, a greater number of links may become established from them, such that activity organisation dominates. Consequently, multiple organisations provide direct access to an event, but its activity provides the best access. Results from Experiment 3 in Conway and Bekerian (1987) argue against this view, because extended events provided better cues than activities.

The experiments reported here investigated the extent to which non-activity event characteristics organise event memories globally and provide direct access to these memories unmeditated by activity. According to the strong activity position, only activities provide direct access to events. According to the weak activity position, multiple organisations provide direct access to events, with activities providing the best access. According to other cross-classification views, multiple organisations provide direct access, with there being no major differences between organisations, or with some non-activity organisation providing the dominant organisation.

The critical measures in these experiments concern the clustering of events during recall. We assume, like many previous investigators, that clustering reflects underlying memory organisation (e.g. Barsalou & Sewell, 1985; Chase & Ericsson, 1981; Graesser & Mandler, 1975; Mandler, 1967; Reitman, 1976; Tulving, 1962, 1964, 1966; Tulving & Pearlstone, 1966). When people attempt to recall information from long-term memory, they access chunks of information sequentially. As people access each chunk, they report as many of its members as are accessible. People do not recall one member from one chunk, one member from a second chunk, one member from a third chunk, return to the first chunk and recall a second member, etc. Instead, people retrieve as many members from the first chunk as possible, proceed to a second chunk, a third chunk, and so forth, rarely returning to a previously accessed chunk.

The logic of our studies rests on these assumptions about memory. If subjects only form global chunks according to activity, then every high-level cluster during recall should share an activity. Within an activity cluster, subclusters may share a non-activity feature (e.g. Events 2 and 5 in Fig. 1 share friends); however, subjects should not form a cluster that does not share an activity, because this would involve retrieving one memory at a time from different clusters (e.g. to cluster memories globally by participant in Fig. 1). In contrast, the organisation in Fig. 2 represents a much wider variety of global chunks, including chunks for participants, locations, and objects. If subjects are organising event memories in this manner, then they could easily produce many non-activity clusters as well as activity clusters.

To implement experimental control over chunking and clustering, we examined people's memory for laboratory events. Most previous studies of event memory have observed people's memory of natural events. However, one problem with natural events is lack of control over their content, distribution, and processing. Laboratory events have obvious advantages in these regards. Of course the problem with laboratory events concerns the generalisability of their results. However, there are two responses to this problem. First, we have performed comparable studies with natural events. If we observe the same pattern of results in both lines of work, then this pattern is unlikely to reflect a particular method. As we shall see in the general discussion, findings from analogous experiments on the organisation of natural events corroborate our findings here for laboratory events (Barsalou, 1988).

Second, there are theoretical reasons for believing that the recall organisation of laboratory events should parallel the recall organisation of natural events. When presented with a set of unfamiliar events to recall, the easiest approach would be to use the same organisational system for recalling natural events. The large literature on organisation in the free recall of taxonomic items supports this assumption. As reviewed in Crowder (1976, Ch.10) and Puff (1979), people use well-established conceptual taxonomies in long-term memory to encode and retrieve lists of words from categories such as *birds*, *furniture*, and *fruit*. When people receive random lists of words from these categories in laboratory settings, they organise them with these well-established taxonomies. If people similarly have well-established organisations for events—which certainly seems reasonable given how central events are to everyday activity—then these organisations should manifest themselves when people encounter laboratory events. Robinson's (1986) finding that people use temporal organisations of the year to learn laboratory lists incidentally indicates that people use event-related organisations to encode laboratory information.

Besides manipulating type of organisation in the experiments to follow, we also manipulated incidental versus intentional learning. Under natural circumstances, people don't try to memorise events as they are experiencing them. Instead, people typically acquire information about events incidentally. In each experiment, therefore, subjects were first exposed to events while performing an incidental orienting task. Following this initial presentation, subjects performed an unexpected recall of the events, enabling observation of organisational tendencies that might occur typically in natural settings when people recall events they did not try to learn. After this initial recall, subjects were told that they would be presented with the same events in a new order and that they would be asked subsequently to remember them again. Of interest was whether different organisational strategies would come into play when subjects tried explicitly to remember the events. This multi-trial recall format enabled us to observe both incidental and intentional memory of laboratory events.

EXPERIMENT 1

The primary purpose of this first study was to explore the extent to which four event characteristics dominate subjects' clustering of laboratory events during free recall: activity, participant, location, and time. Each event contained these four types of information and belonged to an event cluster sharing at least one of these characteristics.

In the low-similarity list, clusters of events shared one and only one of these characteristics (e.g. several events that were instances of *presenting an award*; several other events that occurred in *Hawaii*). If activities provide the dominant organisation of event memories, then we should primarily observe clustering of those events that share an activity—we should not see clustering for events that share only a participant, a location, or a time. If event memories are only stored in organisations for activities, then events that share a participant should be stored in different activity organisations, thereby not being clustered together (e.g. Events 1 and 4 in Fig. 1). Similarly, events sharing only a location or a time should be distributed throughout memory as well. In contrast, if people also organise events with respect to non-activity information, then we should see other kinds of clustering in addition to clustering by activity. For example, people might cluster events that share a participant but not an activity.

In the high-similarity list, each event cluster shared two characteristics (e.g. several instances of *presenting an award* in *Hawaii*). Increasing the similarity of events within a cluster should increase the salience of the clusters and thereby increase clustering at recall. Importantly, however, if activities form the dominant organisation, with participant, location, and time only organising memories subordinately, then this increase in similarity should only occur for clusters that share an activity. Events not sharing an activity should not be clustered together even though they share two characteristics. For example, Events 1 and 4 in Fig. 1 would not be stored together even though they share both a participant (family) and a location (beach). Instead, the increase in similarity should only improve clustering for events that share an activity (e.g. Events 2 and 5, which share eating as well as friends).

In contrast, if people cross-classify event memories in multiple organisations, then we should observe an increase in clustering even for clusters that do not share an activity. Consider Events 1 and 4 in Fig 2, which cluster under global organisations for both participant and location but do not share an activity. Because subjects can retrieve this cluster while searching either of two organisations, they are more likely to retrieve it than to retrieve low similarity clusters, which can only be reached from a single organisation. Thus, the cross-classification view predicts a general advantage for high-similarity clusters over low-similarity clusters, regardless of whether they share an activity.

This experiment also explored the possibility that people pivot between clusters by switching organisations over the course of recall. Consider a possible recall of some fictional events: (1) Paul Newman sorted clothes into light and dark piles; (2) Carl Sagan hung the laundry to dry; (3) Margaret Thatcher loaded the washing machine; (4) Margaret Thatcher painted a picture; (5) Margaret Thatcher pruned the tree. In this sequence, the first three events all share an activity (doing laundry), whereas the last three events all share a participant (Margaret Thatcher). The third event serves as the pivot event, shifting the recall organisation from activity to participant. Pivoting has been observed both in the natural recall of events (Barsalou, 1988) and in children's recall of word lists (Ayres, 1982; Ceci & Howe, 1978; Salatas & Flavell, 1976). In the following experiment, we linked pairs of clusters by pivot events to observe whether subjects would use such pivots to shift between different organisations of laboratory events. Because such pivoting involves event clusters that do not share an activity (e.g. events involving Margaret Thatcher), the strong activity view does not predict it.

Method

Design and Materials. Subjects received the same 36 event descriptions in each of two study periods and attempted to free recall the events after each presentation. Learning was incidental for the first trial and intentional for the second.[2] A low-similarity list and a high-similarity list were constructed to implement the similarity manipulation. Each event in both lists was described by four sentences, with each sentence describing one characteristic of the event (i.e. activity, participant, location, time).

Two versions of each list were constructed. Within each version, events were ordered randomly, with the constraint that events from the same cluster were not adjacent. Across the two versions, the four sentences describing a given event were presented in different random orders, with the constraint that each characteristic occurred equally often in every position. Table 1 provides examples of event clusters from the low-and high-similarity lists.

Low-similarity List. Four different types of clusters were formed, varying in the type of characteristics shared by the events (i.e. activity, participant, location, or time). Three clusters were constructed for each type of characteristic, resulting in a total of 12 clusters. Each cluster contained three events. One "pivot event" in each cluster also belonged to another "linked" cluster (see Table 1). Consequently each list contained 30 critical events, with the 6 pivot events occurring in 2 clusters each and the 24 remaining events occurring in 1 cluster each. In addition, each list included 6 buffer events, 3 at the beginning of the list and 3 at the end.

[2]Note that practice is confounded with the transition from incidental to intentional learning (also in Experiment 2). Thus, conclusions about incidental versus intentional learning must be drawn with care. The lack of intentionality effects in other paradigms, however, suggests that it may not be much of a factor here (e.g. Hasher & Zacks, 1979; Hyde & Jenkins, 1969).

TABLE 1
Examples of Linked Event Clusters Used in Experiment 1

Low-similarity Clusters		*High-similarity Clusters*	
The activity is:	bought a boat.	The activity is:	went swimming.
The participant is:	Debra Winger.	The participant is:	Dolly Parton.
The location is:	Hawaii.	The location is:	Dallas.
The time is:	Veteran's Day.	The time is:	April Fool's Day.
The activity is:	bought a boat.	The activity is:	went swimming.
The participant is:	Alan Alda.	The participant is:	Dolly Parton.
The location is:	Dallas.	The location is:	Italy.
The time is:	Mother's Day.	The time is:	early March.
The activity is:	bought a boat.	The activity is:	went swimming.
The participant is:	Barbara Walters.	The participant is:	Dolly Parton.
The location is:	Mexico.	The location is:	Boston.
The time is:	Fourth of July.	The time is:	Labor Day.
The activity is:	presented an award.	The activity is:	piloted a plane.
The participant is:	Barbara Walters.	The participant is:	Walter Mondale.
The location is:	France.	The location is:	Boston.
The time is:	Easter Sunday.	The time is:	Labor Day.
The activity is:	piloted a plane.	The activity is:	programmed a computer.
The participant is:	Barbara Walters.	The participant is:	Meryl Streep.
The location is:	San Diego.	The location is:	Boston.
The time is:	April Fool's Day.	The time is:	Labor Day.

For the low-similarity clusters, the first three events form an activity cluster, and the last three events form a participant cluster (the third event is the pivot that links the two clusters). For the high-similarity clusters, the first three events form an activity/participant cluster, and the last three events form a location/time cluster (the third event is the pivot that links the two clusters).

High-similarity List. Six different types of clusters were formed, based on the six possible pairings of characteristics (i.e. activity–participant, activity–location, activity–time, participant–location, participant–time, location–time). Two clusters were constructed for each type of characteristic pair, resulting in a total of 12 clusters. Each cluster contained three events. A pivot event in each cluster also belonged to another linked cluster (see Table 1). Consequently each list contained 30 critical events, along with 6 buffer events, 3 at the beginning of the list and 3 at the end.

Subjects and Procedure. Forty undergraduates from Emory University participated for course credit in groups of one to four. Twenty subjects received the low-similarity list, and twenty received the high-similarity list. For each list, half of the subjects received each possible sequence of its two versions across the two study periods.

All instructions and stimuli were presented on individual computer workstations. When ready, subjects received events one at a time for seven seconds each. As the four sentences describing an event appeared simultaneously on the screen, subjects were asked to form an image of the event as it might actually occur. When the event was cleared from the screen, a 7-point scale was displayed, and subjects rated the event for how easy it was to image (i.e. the incidental orienting task). Following presentation of all items, subjects were asked unexpectedly to write down as many full or partial events as possible in whatever order they came to mind. Subjects were allowed as much time as they needed, usually around five minutes. When finished, subjects signalled the workstation and received instructions to study the events again, this time without the rating task and expecting a subsequent recall. Following presentation, subjects performed a second free recall.

Results

Recall. Preliminary analyses assessed various measures of subjects' ability to recall the events. The first analysis assessed subjects' recall without concern for whether the recalled events were correct or incorrect. Because every event recalled by every subject contained at least one characteristic from a presented event, no event was a complete fabrication. Low-similarity subjects recalled a mean of 14.10 events on trial 1 and 21.35 on trial 2, whereas high-similarity subjects recalled 14.75 and 28.30 events on trials 1 and 2, respectively. Recall increased from the first to the second trial [$F(1,32) = 331.08$, $MSE = 4.01$, $P < .001$]. Although similarity had no effect [$F(1,32) < 1$, $MSE = 65.71$], it interacted with trial [$F(1,32) = 4.04$, $MSE = 4.01$, $P < .05$]. High-similarity subjects improved slightly more over trials than low-similarity subjects.

A second analysis assessed how frequently subjects recalled different types of event characteristics. Did subjects vary in how often they recalled activities, participants, locations, and times from the events? The number of times a subject recalled presented event characteristics was scored, independent of whether each recalled characteristic was accompanied by other correct characteristics from an actual event. Subjects hardly ever recalled an event characteristic that was not presented. In the few cases they did, it was always a generalisation of a presented characteristic and was not included in this analysis. If a subject recalled a particular characteristic more times than it actually occurred, each occurrence was counted once (this happened rarely). Table 2 provides the relevant means.

Again frequency of recall increased across trials [$F(1,32) = 301.31$. $MSE = 15.38$, $P < .001$]. Characteristics varied in how well they were recalled [$F(3,96) = 126.65$, $MSE = 14.91$, $P < .001$]. Contrasts found that participants (16.76) and activities (16.06) were recalled equally often, both being recalled more often than locations (11.38), which were recalled more often than times (6.28). Characteristic also interacted with trial [$F(3,96) = 17.58$, $MSE = 2.35$, $P < .001$]. Recall improved less across trials for times than for participants, activities, and locations. Again high similarity recall (13.54) was not reliably

TABLE 2
Mean Frequency of Event Characteristics Recalled in Experiment 1

Trial	Condition: Low-similarity	High-similarity
Participant		
1	12.50	12.65
2	19.95	21.95
Activity		
1	11.15	12.35
2	19.35	21.40
Location		
1	6.00	8.60
2	13.50	17.40
Time		
1	3.35	3.90
2	7.75	10.10

The total number that could have been recalled for each type of characteristic was 36.

better than low similarity recall (11.69) [$F(1,32) = 1.58$, $MSE = 173.50$, $P = .21$]. This general pattern of results suggests that participants and activities were more central to subjects' conceptualisations of the events than were the locations and times in which they occurred.

A third analysis assessed subjects' correct recall of the events. Each recalled event was scored as correct if it contained three of the four characteristics from a presented event. In all cases, three correctly recalled characteristics uniquely specified a presented event. Whereas the average number of total events recalled by subjects was 18.50, the average number of correctly recalled events was 10.64. Caution should be taken when interpreting events that were recalled but not correct by this criterion. Because subjects may have correctly recalled part of an event, but not enough to satisfy the three-characteristic criterion, the event may not constitute an actual intrusion.

Separate ANOVAs were performed for the low-and high-similarity conditions. Table 3 provides the average frequency of total events recalled per subject for each cluster type.[3] In the low-similarity condition, the average number of the nine total events recalled correctly per cluster type increased from the first trial (1.29) to the second trial (3.10) [$F(1,16) = 57.84$, $MSE = 2.27$, $P < .001$]. Cluster type had no effect for low-similarity clusters [$F(3,48) = 1.56$, $MSE = 0.97$,

[3]Recall that there were three clusters for each of the four types in the low-similarity condition versus two clusters for each of the six types in the high-similarity condition, with every cluster containing three events. Thus, there were 36 critical events that subjects could have recalled in each similarity condition. Because 6 were pivot events, each counting once for each of two clusters, there were actually only 30 critical events in each list.

TABLE 3
Average Total Frequency of Correctly Recalled Events for Each Cluster Type in Experiment 1

Cluster Type	*Trial 1*	*Trial 2*
Low-similarity Condition		
Participant	1.20	2.95
Activity	1.15	2.85
Location	1.35	3.35
Time	1.45	3.25
High-similarity Condition		
Participant/Activity	1.95	3.90
Participant/Location	1.20	2.70
Participant/Time	1.20	2.55
Activity/Location	1.60	2.50
Activity/Time	1.10	2.75
Location/Time	0.90	2.65

An event was coded as correctly recalled if a subject recalled at least three of its presented characteristics. Subjects received three clusters of each type in the low-similarity condition and two clusters of each type in the high-similarity condition, with each cluster containing three events. Thus, the maximum recall for each average was nine events in the low-similarity condition and six events in the high-similarity condition.

$P = .21$]. In the high-similarity condition, the average number of the six total events recalled correctly per cluster type increased from the first trial (1.32) to the second trial (2.84) [$F(1,16) = 72.04$, $MSE = 1.79$, $P < .001$] and varied across cluster type [$F(5,80) = 5.42$, $MSE = 1.30$, $P < .01$].

Clustering. The primary analyses to follow assessed the organisation of subjects' protocols. A cluster was defined as any contiguous sequence of two or three recalled events that shared a common value for at least one event characteristic, such as:

Brooke Shields threw a party in San Diego on St. Patrick's Day.
Brooke Shields threw a party in France.

or:

Geraldine Ferraro went on a diet,
a vacation
and sold a farm.

Note that subjects often omitted event characteristics, sometimes because they could not remember them (e.g. the time for the second Shields event), and sometimes because of ellipsis (the participant in the second and third Ferraro events). When ellipsis was obvious, subjects were credited with recalling the ellipsed information.

Cluster analyses were performed on all recalled events, regardless of whether they were scored as correct or not (as described earlier). We were primarily interested in how subjects organised whatever events they thought they had seen, and the presence of imperfectly recalled events does not interfere with observing such organisation. If we had eliminated incorrect events from these analyses, we would have created gaps in subjects' recall sequences that would have been difficult to interpret. By including all events recalled, we maintained the overall conceptual organisation of each subject's recall. Furthermore, these clusters are also based on inferred event characteristics. For example, if the location for an event was not mentioned, but if the event's other characteristics uniquely identified it as belonging to a particular location cluster, the event was counted as belonging to it. Finally, when buffer events were recalled, they never counted towards a cluster, given each had no shared event characteristics, but they could break up events from a cluster surrounding them. For example, if a buffer event occurred between two events from a location cluster, these two events were not counted as a cluster.

Low-similarity subjects formed an average of 0.90 clusters on trial 1 and 2.30 clusters on trial 2, whereas high-similarity subjects formed an average of 2.50 clusters on trial 1 and 5.00 clusters on trial 2.[4] As predicted, more clustering occurred for high-similarity subjects (3.75) than for low-similarity subjects (1.60) [$F(1,38) = 12.54$, $MSE = 7.37$, $P < .01$]. Clustering increased from trial 1 (1.70) to trial 2 (3.65) [$F(1,38) = 41.94$, $MSE = 1.81$, $P < .001$], with similarity and trial interacting marginally [$F(1,38) = 3.34$, $MSE = 1.81$, $P < .10$]. The increase in clustering across trials was greater for high-similarity subjects than for low-similarity subjects.

The next analysis assessed the relative rate at which subjects produced clusters of different types. The strong activity view predicts that subjects should never have formed a non-activity cluster. The weak activity view acknowledges the production of non-activity clusters but predicts that activity clusters should occur most frequently. Other cross-classification views predict either that no organisation should dominate, or that some non-activity organisation could dominate. Table 4 presents the average frequency of clusters that a subject produced for each organisational type. In the low-similarity condition, cluster type varied, with subjects producing more activity and participant clusters than location and time clusters [$F(3,57) = 4.51$, $MSE = 0.34$, $P < .01$]. Activity clustering was not significantly higher than participant clustering on either

[4] On a given trial, a subject could have formed a maximum of 12 clusters, 3 for each of 4 organisations in the low-similarity condition, or 2 for each of 6 organisations in the high-similarity condition.

TABLE 4
Mean Frequency (ARC Score) of Clusters in Experiment 1

Trial		*Organisation* *Participant*	*Activity*	*Location*	*Time*
Low-similarity Condition					
1		0.20 (0.07)	0.50 (0.05)	0.15 (0.08)	0.05 (0.08)
2		0.70 (0.23)	0.80 (0.27)	0.45 (0.02)	0.35 (0.06)

Trial	*Participant/ Activity*	*Participant/ Location*	*Participant/ Time*	*Activity/ Location*	*Activity/ Time*	*Location/ Time*
High-similarity Condition						
1	0.80 (0.51)	0.45 (0.36)	0.25 (0.21)	0.50 (0.26)	0.45 (0.25)	0.05 (0.00)
2	1.55 (0.80)	0.85 (0.46)	0.80 (0.44)	0.65 (0.40)	0.75 (0.30)	0.40 (0.28)

Trial	*Collapsed Organisation*[a] *Participant*	*Activity*	*Location*	*Time*
1	1.50 (0.36)	1.75 (0.34)	1.00 (0.21)	0.75 (0.15)
2	3.20 (0.57)	2.95 (0.50)	1.90 (0.38)	1.95 (0.34)

The maximum frequency for mean cluster frequency is 3 in the low similarity condition and 2 in the high similarity condition.

[a]See the text for a description of how organisations were collapsed.

trial. Furthermore, cluster type did not interact with trial [$F(3,57) = .33$, $MSE = 0.30$]. These results reject the strong activity view, given that subjects produced non-activity clusters. These results also reject the weak activity view, given that activity clustering did not dominate all other types (i.e. participant clustering was equally frequent).

The results for the high-similarity condition in Table 4 yield similar conclusions. Again, cluster type varied, with participant–activity clusters being most common, location–time clusters being least common, and the other four cluster types being intermediate [$F(5,95) = 10.21$, $MSE = 0.37$, $P < .001$]. Again, cluster type did not interact with trial [$F(5,95) = 1.62$, $MSE = 0.27$]. As in the low-similarity condition, subjects formed clusters that did not share activities, contrary to the strong activity view.

To estimate the relative importance of each organisational type in the high similarity condition the three frequencies for all clusters involving a characteristic were summed. For example, the sum of the participant–activity, participant–location, and participant–time frequencies for trial 1 represents the relative importance of participant organisation on trial 1. These estimates are shown at the bottom of Table 4. Because each uncollapsed mean contributes to two collapsed means, the lack of independence between the collapsed means makes performing statistical tests questionable. However, an examination of the collapsed means suggests that activity and participant organisation were of roughly equal importance on both trials and more important than location and time organisation. This result disconfirms the weak activity view, which assumes that activity organisation should dominate all other types.

As we have seen, subjects formed non-activity clusters. The following analysis explores this finding further. Recall that subjects formed more clusters in the high-similarity condition than in the low-similarity condition. According to the strong activity view, this similarity effect should only involve activity clusters. As described in the introduction to this experiment, increasing the similarity of the events in a cluster should only increase clustering when these events already share an activity. Thus, adding a common participant to several events that share a location but not an activity should not increase the likelihood of clustering them at recall, because they are distributed across different activity organisations. To assess this hypothesis, we computed the average proportion of the possible non-activity clusters produced per subject (i.e. clusters sharing only participant, location, and/or time). Proportions were used as the dependent measure, because the number of possible non-activity clusters differed in the low-(9) and high-(6) similarity conditions.

The probability of forming a non-activity clustering increased from trial 1 to trial 2 [$F(1,38) = 27.44$, $MSE = 0.09$ arcsin units, $P < .001$]. Low-similarity subjects formed 0.04 of the possible non-activity clusters on trial 1 and 0.17 of the possible non-activity clusters on trial 2. The proportions for high-similarity subjects were 0.12 on trial 1 and 0.34 on trial 2. Most importantly, the overall probability of forming a non-activity cluster was higher for high-similarity subjects (.23) than for low-similarity subjects (.11), indicating that non-activity clusters contributed to the similarity effect [$F(1,38) = 10.72$, $MSE = 0.20$ arcsin units, $P < .01$]. Contrary to the strong activity view, higher event similarity increased clustering, even when the events in a cluster did not share an activity. Trial and cluster type did not interact [$F(1,38) = 1.70$, $MSE = 0.09$ arcsin units].

ARC Analysis. The average frequency of clusters in Table 4 is not corrected for differential recall of events from the various cluster types. To the extent that some cluster types are better recalled than others, they have a greater probability of producing clusters by chance. To correct for this possibility, the Associated Ratio of Clustering (ARC) was applied to the protocols (Roenker, Thompson, & Brown, 1971). When clustering occurs at chance levels, ARC scores approximate 0; when clustering is perfect, ARC scores approximate 1. Table 4 presents ARC scores by condition in parentheses.

In the low-similarity condition, there was no overall effect of cluster type [$F(3,57) = 1.14$, $MSE = 0.096$]. Clustering increased from trial 1 to trial 2 (0.07 vs. 0.14), but this difference only approached significance [$F(1,19) = 2.14$, $MSE = 0.105$, $P = .16$]. Cluster type and trial did not interact [$F(3,57) = 1.48$, $MSE = 0.131$]. On trial 1, clustering was significantly greater than zero [$t(19) = 2.09$, $SE = 0.033$, $P < .05$], but cluster types did not differ.

On trial 2, clusters for participants and activities were marginally more frequent than clusters for locations and times. Notably, activity clustering failed to dominate all other forms of clustering, contrary to the strong and weak activity views.

Clustering increased reliably from the low-similarity condition (0.11) to the high-similarity condition (0.36) [$F(1,38) = 18.26$, $MSE = 0.07$, $P < .001$]. Within the high-similarity condition, clustering increased from trial 1 (0.26) to trial 2 (0.45) [$F(1,19) = 19.62$, $MSE = 0.10$, $P < .001$] and varied across cluster types [$F(5,95) = 5.43$, $MSE = 0.21$, $P < .001$]. Trial and cluster type did not interact [$F(5,95) = .52$, $MSE = 0.18$]. Participant–activity clustering was most prevalent, followed by clusters that shared either participants or locations. When clustering principles were collapsed (Table 4), activity clustering failed to dominate all other forms of clustering, being slightly less than participant clustering.

Together, the results for cluster frequency and ARC scores cast doubt on the strong and weak activity views. In contrast to the strong activity view, subjects formed clusters of events that did not share activities. In contrast to the weak activity view, activity clusters did not dominate all other types, given that participant clusters occurred as frequently as activity clusters. In general, these results support the conclusion that subjects cross-classified events into multiple organisations at the global level.

Cluster Length. The mean length (i.e. number of events) per cluster was computed for each characteristic in the low-similarity condition and each pair of characteristics in the high-similarity condition. Each mean length was computed across only those subjects producing clusters of that type. High-similarity subjects produced longer clusters than low-similarity subjects on both trial 1 [2.24 vs. 2.08; $t(64) = 8.43$, $SE = 0.07$, $P < .001$] and trial 2 [2.35 vs. 2.18; $t(138) = 2.36$, $SE = 0.08$, $P < .05$]. Clusters were longer on trial 2 than on trial 1 in the high-similarity condition [$t(143) = 3.96$, $SE = 0.07$, $P < .001$] but not in the low-similarity condition [$t(59) = 1.21$, $SE = 0.10$]. Activity–participant and activity–location clusters were both longer than activity–time clusters [$t(69) = 2.36$, $SE = 0.12$, $P < .05$; $t(45) = 2.36$, $SE = 0.12$, $P < .05$]. Other than these effects, there was little variation in cluster length across cluster types. Lancaster (1985) provides further details about the findings on cluster length.

Pivoting. Pivot events were defined as a recalled event that ended one cluster in a subject's protocol but simultaneously began another cluster. For example, the following protocol segment illustrates a pivot from a participant–location cluster to an activity–location cluster:

Geraldine Ferraro programmed a computer in New Orleans.
Geraldine Ferraro started a diet in New Orleans.
Woody Allen started a diet in Vermont.
Dolly Parton started a diet in Vermont.

TABLE 5
Type (Frequency) of Pivots in Experiment 1

	Condition	
Trial	*Low-similarity*	*High-similarity*
1	P → A (1)	PA → L (1)
	A → P (1)	PL → AL (1)
		P → AL (1)
2	A → P (2)	LT → PA (2)
	P → L (1)	P → A (1)
		L → PA (1)
		A → PL (1)
		AL → T (1)
		PL → A (1)
		PL → T (1)

x → *y* means that a subject pivoted from a cluster of type *x* to a cluster of type *y*, where P, A, L, and T refer to participant, activity, location, and time, respectively.

Table 5 shows the number and types of pivots that subjects produced. The characteristics defining each pivot are indicated by the single letters, A, P, L, and T, which represent activity, participant, location, and time clusters, respectively. The arrows represent the transition from the first cluster in each pivot to the second. Note that high-similarity subjects sometimes constructed clusters based on one characteristic instead of two. Pivoting occurred more frequently on trial 2 (11 occurrences) than on trial 1 (5 occurrences).

The pattern across both similarity conditions was for pivoting to occur only in the presence of clustering by participant or activity, given that 100% of the pivots involved either or both of these characteristics. Overall, neither participant or activity dominated, with 94% of the pivots involving participant and 88% involving activity. Across pivots, 56% were a transition from a cluster sharing at least a participant or activity to a cluster sharing at least an activity or participant; 31% were the result of pivoting from a cluster sharing at least a participant or activity to a cluster sharing at least location.

Discussion

These results contradict the strong activity view of event organisation, which proposes that subjects should never cluster events that do not share an activity. On the contrary, subjects formed clusters that did not share an activity but that shared a participant, location, or time. Furthermore, non-activity clusters contributed to the similarity effect in clustering. As the similarity of events within a cluster increased, the probability of clustering them at recall increased as well, even for clusters that did not share an activity. Finally, subjects pivoted between different organisations, switching to and from non-activity clusters. Together, these three findings indicate that non-activity forms of organisation are not embedded in larger clusters of events that share an activity (Fig. 1). Instead, people organise events into clusters that do not share an activity but that share other event characteristics, such as participants, locations, and times (Fig. 2).

Although subjects formed clusters around locations and times, they tended to prefer clusters for participants and activities. Interestingly, subjects formed participant clusters as often as activity clusters. On incidental trial 1 in both the low-and high-similarity condition, subjects were just as likely to cluster by participant as by activity. On intentional trial 2 in both similarity conditions, participant clusters were just as likely as activity clusters. Furthermore subjects recalled participants as highly as activities in all conditions; the length of participant clusters was comparable to the length of activity clusters; and participants were as central to pivoting as activities. Together, all of these results suggest that participants are no less central to event organisation than activities. At the least, these results indicate that activities do not constitute the sole organisation of event memories. People do indeed organise event memories according to other event characteristics as well.

EXPERIMENT 2

Activities and participants were equally central to subjects' organisations of event memories in Experiment 1. To better assess the weak activity view—that activities dominate all other organisations—Experiment 2 examined activity and participant organisations more closely. In this experiment, every event could be clustered into either an activity cluster or a participant cluster but not both, thereby pitting activity and participant organisation against each other. At one extreme, subjects could organise the entire list into activity clusters; at the other extreme, subjects could organise the entire list into participant clusters. Intermediately, subjects could integrate some events into activity clusters and others into participant clusters. Similar to Experiment 1, the potential for pivoting between activity and participant clusters was built into the event structure. Because every event in an activity cluster also belonged to a participant cluster, and vice versa, every critical event in the experiment constituted a potential pivot between two clusters. Unlike Experiment 1, cluster size varied from two to four, thereby preventing subjects from always assuming a cluster size of three. We also increased the number of study–test trials from two to five, thereby enabling the observation of increasing organisation over time. Similar to Experiment 1, trial 1 was incidental, and trials 2–5 were intentional.

Experiment 2 was also designed to assess individual differences. If we observe a mixture of activity and participant clustering, how are these two forms of organisation distributed within subjects? Does every subject adopt only one type of organisation, with some subjects forming only activity clusters and other subjects forming only participant clusters? Or does a given subject often mix activity and participant clusters, perhaps linking them together through pivots? Furthermore, does a given subject maintain a constant form of organisation across study-recall trials? Or do subjects change their organisation?

The final purpose of Experiment 2 was to compare three possible types of activity clusters. In identity clusters, each event in a cluster shared an identical activity (e.g. *rode the bus*). In superordinate clusters, each event in a cluster contained a different activity that was an instance of the same superordinate activity (e.g. *had the car tuned* and *had the brakes adjusted* as instances of *car repairs*). In script clusters, each event in a cluster contained an activity that was a part of the same script (e.g. *read the menu* and *ordered the daily special* from the *restaurant* script). Of interest was whether subjects would be more likely to organise events for one particular activity organisation over another.

Method

Design and Materials. Subjects received the same 48 event descriptions on each of five study trials and attempted to free recall the events after each presentation. Learning was incidental for the first trial and intentional for the remaining four trials. Subjects received different random orders of the events on a given trial, and a given subject received different random orders across trials.

A single sentence described each event, providing both the participant and activity. Each of 12 event clusters was defined by a shared activity. Identical, superordinate, and script clusters (as just described) each constituted one third of the clusters. For each cluster type, one cluster contained two events, two contained three events, and one contained four events, for an average of three events in each of 12 clusters (36 events total). An additional 12 events were included whose activity was unrelated to the activity of any other event. These events served the purpose of diluting the organisational density of the list.

Orthogonal subsets of the 36 events that formed activity clusters also shared participants. Three of these clusters contained two events, six contained three events, and three contained four events (for the same distribution of cluster sizes as the activity condition). The 12 events that did not share an activity with any other event also did not share a participant with any other event. The participant clusters were uncorrelated with the activity clusters in that no two events in the same activity cluster ever shared the same participant, and no two events within the same participant cluster ever shared a related activity. Every event in a cluster provided a potential pivot between an activity and a participant cluster, because each of these events belonged simultaneously to one cluster of each type. Table 6 provides examples of the event clusters, with examples of participant clusters being distributed across the blocked activity clusters.

TABLE 6
Examples of Event Clusters Used in Experiment 2

	Event Characteristic	
Activity Type	*Participant*	*Activity*
Identity	Jane Fonda	rode the bus to Macon.
	John Belushi	rode the bus to Columbus.
	Barbra Streisand	rode the bus to Athens.
Superordinate	Ryan O'Neill	had the car tuned.
	John Belushi	had the brakes adjusted.
	Barbra Streisand	had the transmission rebuilt.
Script	Robert Redford	read a menu.
	Bo Derek	ordered the daily special.
	Barbra Streisand	ate an enjoyable meal.
Unrelated	Liz Taylor	wrote a letter.
	Walter Cronkite	entered a sweepstakes.
	Brooke Shields	knit a sweater.

Examples of participant clusters are shown for John Belushi and Barbra Streisand.

Two versions of the list were constructed. In one, events were described in active voice sentences so that the participant was named before the activity (e.g. Liz Taylor wrote a letter). In the other, events were described in passive voice sentences so that the activity preceded the participant (e.g. A letter was written by Liz Taylor). Both active and passive voice lists were used to control for any possible organisational bias that might result from the order in which subjects read the participants and activities.

Subjects and Procedure. Twenty-four undergraduates from Emory University participated in groups of one to four. Twelve subjects received the stimulus list in the active voice, and twelve received it in the passive voice.

All instructions and stimuli were presented on individual computer workstations. When ready, subjects received events one at a time for seven seconds in a random order. After each sentence was cleared from the screen, a rating scale was displayed, and subjects rated the event for its interestingness (i.e. the incidental orienting task). Following presentation of all items, subjects were asked unexpectedly to write down as many full or partial events as possible, in whatever order they came to mind. Subjects were allowed as much time as they needed, usually around five minutes. When finished, subjects signalled the workstation and received instructions to study the events again, this time without the rating task and expecting a subsequent recall. Following presentation of the events in a new random order, subjects performed another free recall. Subjects continued the study–recall cycle described for the second trial three more times for a total of five trials.

Results

At each step of analysis, the data were first checked for differences between the active and passive voice versions of the list. Because no differences occurred, all results are reported collapsed across this manipulation.

Recall. The recall data were scored for the number of the 48 events recalled, regardless of whether the events were complete or correct. Column 2 of Table 7 provides the relevant means. Recall increased across trials $F(4,92) = 35.03$, $MSE = 47.21$, $P < .001$]. The items were then coded for completeness and correctness. An item was considered complete if both a participant and activity were recalled. An item was considered correct if the participant and activity belonged to the same stimulus event. Relevant means are shown in columns 3 and 4 of Table 7. Recall increased across trials for both complete and correct events [$F(4,92) = 46.29$, $MSE = 45.73$, $P < .001$; $F(4,92) = 49.26$, $MSE = 44.53$, $P < .001$].

Overall the probability of correct recall was higher for participants (.60) than for activities (.55) [$F(1,23) = 20.01$, $P < .001$]. However, organisation interacted with trial [$F(4,92) = 10.01$, $P < .001$], with the difference between recall for participants and activities only being higher on trial 1 (0.36 vs. 0.25) and not differing thereafter (0.66 vs. 0.63 for trials 2–5). Within activities, the total number of events recalled correctly differed for each of the three activity types [$F(2,46) = 4.56$, $MSE = 7.63$, $P < .05$], with recall being best for the identity clusters (7.80), then the superordinate clusters (7.20), and finally the script clusters (6.72). Decreasing similarity within clusters may have produced this effect.

The probability of accessing a cluster increased with its size [$F(2,44) = 28.50$, $MSE = 0.07$, $P < .001$], where access was defined as retrieving one or more events from a cluster (Tulving & Pearlstone, 1966). For activity clusters, the probability of access was .66 for two-event clusters, .81 for three-event clusters,

TABLE 7
Mean Number of Events Recalled in Experiment 2

		*Criterion**	
Trial	*Total*	*Complete*	*Correct*
1	17.67	11.92	10.71
2	26.71	23.50	22.33
3	33.71	31.25	30.12
4	36.45	33.33	32.62
5	37.62	34.42	33.58

*See the text for descriptions of the criteria. Forty-eight possible events could have been recalled on each trial.

and .83 for four-event clusters (averaged across trials). For participant clusters, the probability of access was .69 for two-event clusters, .86 for three-event clusters, and .83 for four-event clusters (averaged across trials). In contrast, the probability of retrieving a given item decreased with cluster size [$F(2,44) = 13.50$, $MSE = 0.05$, $P < .001$], namely, output interference (e.g. Neely, Schmidt, & Roediger, 1983; Roediger, Neely, & Bladton, 1983). For activity clusters, the probability of retrieving a given item was .74 for two-event clusters, .71 for three-event clusters, and .69 for four-event clusters (averaged across trials). For participant clusters, the probability of retrieving a given event was .75 for two-event clusters, .70 for three-event clusters, and .60 for four-event clusters (averaged across trials). Lancaster (1985) provides further details about the effects of cluster size and its interactions with other variables.

Clustering. The number and length of clusters produced by each subject were computed for each activity type, for all activity types combined, and for participants. If a subject recalled two or more events from an *a priori* event cluster, the subject was credited with forming a cluster for it. Table 8 shows the relevant means.

Number of Clusters. In an analysis comparing clustering within the three activity types (columns 2, 3, and 4 in Table 8), number of clusters increased over trials [$F(4,92) = 21.95$, $MSE = 1.21$, $P < .001$]. Activity type also differed, with identity clusters, not surprisingly, producing more clustering than script and taxonomic clusters [$F(2,46) = 24.82$, $MSE = 1.12$, $P < .001$]. Trial and activity type did not interact [$F(8,184) = 1.87$, $MSE = 0.54$].

In a second analysis comparing overall activity clustering with clustering by participant (columns 5 and 6 in Table 8), clustering again increased across trials [$F(4,92) = 25.20$, $MSE = 5.56$, $P < .001$]. Although activity clusters generally

TABLE 8
Mean Frequency (ARC Score) of Clusters in Experiment 2

	Cluster Type				
Trial	*Script*	*Superordinate*	*Identity*	*Activity*[a]	*Participant*
1	0.46 (0.31)	0.54 (0.13)	0.92 (0.41)	1.92 (0.33)	1.71 (0.29)
2	0.96 (0.41)	1.33 (0.58)	2.29 (0.69)	4.58 (0.65)	3.29 (0.40)
3	1.62 (0.52)	1.92 (0.53)	2.50 (0.82)	6.04 (0.66)	4.87 (0.41)
4	1.75 (0.41)	1.58 (0.41)	2.46 (0.67)	5.79 (0.54)	5.67 (0.42)
5	1.83 (0.47)	1.71 (0.55)	2.83 (0.87)	6.37 (0.64)	5.25 (0.38

The total sum of participant and activity clusters that could have been recalled on each trial was 12. Presented clusters ranged in length from two to four events, with an average of three.

[a]Frequencies in this column are sums across the three activity types in the first three columns; ARC scores were computed separately for each of the three activities and for all activities together.

occurred slightly more frequently than participant clusters, they did not differ reliably [$F(1,23) < 1$, $MSE = 37.33$], nor did they interact [$F(4,92) < 1$, $MSE = 4.66$].

ARC Analysis. To control for different cluster frequencies in different conditions, ARC analyses were performed on subjects' protocols. Table 8 presents the average ARC score for each condition in parentheses. In general, the results are the same as those for the analysis of cluster frequency. In an analysis comparing the three activity types and participants, clustering increased across trials [$F(4,92) = 11.07$, $P < .001$], and with increasing similarity of the different organisations [$F(3,69) = 5.92$, $P < .005$].

In a second analysis comparing overall activity clustering with clustering by participant, clustering again increased across trials [$F(4,92) = 6.60$, $P < .001$]. Although activity clustering was generally higher than participant clustering, this difference only approached significance $F(1,23) = 2.77$, $P = .10$]. Notably, activity and participant clustering did not differ on trial 1 under incidental learning.

Cluster Length. Because not all subjects produced clusters of each type on each trial, *t*-tests were performed on only those subjects who produced clusters in each cell of the design. Given the small number of clusters for each activity type, analyses of length were performed on only the three activity types combined in comparison to participants. On trial 1, activity and participant clusters did not differ in length [2.07 vs. 2.12; $t(36) < 1$, $SE = 0.08$]. On trial 2, activity and participant clusters differed reliably [2.52 vs. 2.15; $t(39) = 3.51$, $SE = 0.10$, $P < .01$] and remained reliably different on trial 3 [2.57 vs 2.26; $t(43) = 3.36$, $SE = 0.09$, $P < .01$], trial 4 [2.66 vs. 2.25; $t(42) = 4.12$, $SE = 0.10$, $P < .01$] and trial 5 (2.61 vs. 2.26; $t(42 = 4.29$, $SE = 0.11$, $P < .01$]. The only effect between adjacent trials was that activity clusters were reliably longer on trial 2 than on trial 1 [$t(36) = 4.50$, $SE = 0.10$, $P < .01$].

Pivoting. Subjects often combined individual clusters that were based on different organising principles into coherent strings of clusters by pivoting from one organising principle to the other around a single event (as described in Experiment 1). Pivoting was made possible in this experiment by every critical event belonging to both an activity and a participant cluster (see Table 6 for examples). A *string* was defined as a series of two or more linked clusters, where every cluster (except the last) was linked to the following cluster by the recall of a pivot event. A string was coded as either an activity or participant string, depending on whether the cluster that began it was an activity or participant cluster. Table 9 provides the mean number of activity and participant strings for strings having two clusters and for strings having three or more clusters.

TABLE 9
Mean Frequency of Strings Produced per Subject and Frequency of Subjects Producing Strings (in Parentheses) in Experiment 2

Trial	Type of String: Participant	Type of String: Activity
Strings of Two Clusters		
1	0.12 (4)	0.21 (5)
2	0.33 (6)	0.54 (10)
3	0.37 (8)	0.79 (12)
4	0.79 (12)	0.46 (9)
5	0.62 (11)	0.75 (12)
Strings of Three or More Clusters		
1	0.04 (1)	0.04 (1)
2	0.12 (3)	0.25 (5)
3	0.42 (8)	0.29 (5)
4	0.42 (8)	0.29 (7)
5	0.29 (7)	0.29 (6)

The number of subjects who could have produced strings in each condition was 24.

The number of strings increased across trials [$F(4,92) = 7.87$, $MSE = 0.34$, $P < .001$], with two-cluster strings being more common than longer strings [$F(1,23) = 25.12$, $MSE = 0.31$, $P < .001$]. Activity and participant strings did not differ [$F(1,23) < 1$, $MSE = 0.40$], and no interactions were significant.

Table 9 also shows in parentheses the number of subjects who produced strings of various types on a given trial. A given subject may be represented in more than one cell on any given trial if that subject's protocol contained strings of different types and/or lengths. As can be seen from Table 9, the number of subjects producing each type and length of string roughly paralleled the results for number of strings.

Individual Differences. The analyses presented so far do not differentiate between the possibility that all subjects used both organising principles versus the possibility that subjects differed systematically in their preferred organisation. To explore this issue, we assessed three individual difference measures. *Total clusters* measured the number of activity clusters produced by a subject, divided by his or her total number of clusters. *Strings* measured the number of activity strings, divided by the total number of strings. *Independent clusters* measured the number of activity clusters that did not belong to a string, divided by the total number of clusters not belonging to a string.

Scores ranged from 0 to 1 for each measure, with values near 0 representing a preference for participant organisation and values near 1 representing a preference for activity organisation. Subjects with scores of .33 or lower were considered to be ‘‘participant dominant’’, subjects with scores between .33 and .67 were considered to be ‘‘equivalent’’, and subjects with scores of .67 or higher were considered to be ‘‘activity dominant’’. Table 10 provides the distribution of subjects across trials for each measure.

On trial 1, a majority of subjects were activity dominant on all three measures. However, a quarter of the subjects were participant dominant on trial 1, as measured by total clusters, and another quarter used participant and activity clustering equivalently. Furthermore, the total clusters results show that subjects evolved towards being equivalent on trials 2 through 5. By trial 5, roughly equal numbers of subjects were activity dominant and equivalent for this measure. In contrast, the strings measure remained polarised, given that few subjects began strings with activity and participant clusters equivalently. Instead, the percentage of subjects beginning strings with activity clusters hovered around 60%, whereas the percentage of subjects beginning strings with participant clusters increased from 17% to 41%. Finally, subjects’ preferences for independent clusters remained fairly constant over trials, with about twice as many subjects always preferring independent activity clusters as subjects preferring independent participant clusters.

Most subjects remained relatively consistent in their organisation across trials. Of the 24 subjects, only 6 showed any major change in their preferred

TABLE 10
Proportions of Subjects Exhibiting Particular Organisational Styles

		Organisational Style	
Trial	*Participant Dominant*	*Equivalent*	*Activity Dominant*
Total Clusters			
1	.25	.21	.54
2	.12	.50	.38
3	.17	.46	.38
4	.25	.50	.25
5	.12	.42	.46
Strings			
1	.17	.17	.67
2	.13	.13	.73
3	.33	.17	.50
4	.62	.10	.29
5	.41	.00	.59
Independent Clusters			
1	.29	.10	.62
2	.22	.22	.56
3	.29	.12	.60
4	.30	.13	.56
5	.35	.09	.56

mode of organisation. Each of these six subjects shifted on trial 2 from an activity dominant organisation to an organisation that utilised activity and participant organisations equivalently. All but one of the remaining 18 subjects (who could not be classified definitively overall) maintained one of the three organisational styles throughout the five trials. Of these subjects, 13 were consistently activity dominant, 3 were consistently participant dominant, and 1 was consistently equivalent.

Discussion

Similar to the results of Experiment 1, the results of Experiment 2 disconfirm the strong activity view. Subjects recalled participants as highly as activities. Although subjects formed slightly more activity clusters than participant clusters, this difference was not reliable across trials. Nor was it reliable on any given trial, including the first incidental trial and the later intentional trials. At the level of individual subjects, many subjects formed participant clusters from the first trial on, with some subjects strongly preferring participant clusters. Across the four intentional trials, subjects became increasingly likely to use participant organisation at the expense of activity organisation. Together, all of these results indicate that subjects readily form clusters of events that do not share activities.

However, several findings support the weak activity view. Although subjects frequently formed participant clusters, they tended to form activity clusters more often, and activity clusters also tended to be longer. At the level of individual subjects, more subjects remained activity dominant over the course of the experiment than ever became participant dominant. In these regards, activity organisation exhibited a modest advantage over participant organisation.

Nevertheless, many subjects used both organisations, even when they could have used one or the other exclusively to organise all of the potentially organisable events. On trial 1, 21% of the subjects used activity and participant organisation equivalently. By trial 5, 42% of the subjects were using these two organisations equivalently. These results demonstrate clearly that people can simultaneously organise events using two forms of global organisation. People are not limited to organising events only by activity at the global level. Instead, people also organise events simultaneously into multiple organisations, as cross-classification views predict. The extensive pivoting that produced strings further supports cross-classification views, because pivoting can only occur if at least some events are cross-classified in multiple organisations.

GENERAL DISCUSSION

The findings from these studies cast doubt on the strong activity view of event organisation. People do not organise events globally only by activities, with other forms of organisation embedded locally within activity clusters (Fig. 1). Instead, people cross-classify events into multiple global organisations simultaneously (Fig. 2). In both Experiments 1 and 2, subjects formed global clusters of events that did not share activities but instead shared participants, locations, and/or times. Furthermore, clustering in Experiment 1 increased when the number of shared non-activity characteristics increased from one to two. Even though these events shared no activities, subjects were still able to recognise and utilise the increase in non-activity similarity to achieve greater organisation. Finally, subjects pivoted between activity and non-activity clusters in both experiments, further indicating that they had cross-classified events in multiple organisations. Once subjects exhausted a cluster of events that shared a characteristic (e.g. swimming), they were able to focus on a different characteristic of the last event recalled (e.g. Walter Cronkite) and shift the focus of search to events sharing it, even when these events did not share an activity.

Nevertheless, some of the findings from Experiment 2 supported the weak activity view of organisation. Although these subjects tended to cross-classify events in multiple organisations globally, they exhibited a modest preference for activity organisation. The dominance of activities could reflect their importance in understanding events. During the comprehension of an event, people must explain many of its aspects with respect to the goals and plans that underlie the basic activity. As a result, many relations develop between the event and activity knowledge. To the extent that more relations develop to activity knowledge than to other organisational knowledge, the activity may provide a stronger organisational cue (Barsalou, 1995).

However, caution is necessary in drawing conclusions about the dominance of any single organisation. In Experiment 1, subjects used participants as often as activities to cluster events. In Experiment 2, participant organisation was comparable to activity organisation on some key measures. Finally, investigators have observed other dominant forms of organisation on occasion. For example, Conway and Bekerian (1987) and Barsalou (1988) found that extended events were more important organisers of events than activities. Brown et al. (1986) similarly observed the importance of extended events. Although activities are certainly important organisers, they may not always be most important. Further comparisons of activities and extended events is necessary, and one useful tack might be to compare them in memory for laboratory events.

A second reason for caution concerns the role of context. Depending on the goals of the perceiver, different organisations may become salient. Activities are important for strategists, locations are important for geographers, participants are important for biographers, times are important for historians, and so forth. As people pursue different goals, different organisations may become salient and necessary for explaining relevant events. Nevertheless, certain organisations may be of general importance across goals, such as activities, with other organisations becoming more or less important depending on context. Furthermore, extensive cross-classification may always occur, with events being organised into multiple organisations simultaneously but with their relative dominance varying. To the extent that one organisation receives extensive processing, such that many relations develop to event memories, it will dominate organisation.

Caution is also warranted in generalising results from studies of laboratory events to conclusions concerning natural events. However, two factors suggest that such conclusions have some basis. First, similar results have been obtained in naturalistic settings. Barsalou (1988) found that people used participants, locations, and times—not just activities—to organise autobiographical events globally. In a series of studies that assessed the free recall, cued recall, and timed retrieval of autobiographical memories, activity cues never dominated. Indeed, temporal, participant, and location cues dominated activity cues on occasion. Barsalou (1988) also observed pivoting, which again implies the cross-classification of events in multiple organisations. Thus, the presence of analogous results in autobiographical memory corroborates the results we have obtained in a laboratory setting. In both cases, we have observed evidence that multiple organisations cross-classify event memories.

A second reason for believing that our laboratory results have generality comes from much previous research on taxonomic organisation in memory. Many previous researchers have found that subjects organise lists of words into the familiar taxonomic categories that pervade daily activity. As these results illustrate, people use their natural organisational tendencies to structure information in laboratory situations. Although the laboratory events that subjects received in our experiments are unlike natural events in many ways, it is not unreasonable to expect the organisation of such events to parallel the organisation of natural events. When presented with a set of unfamiliar events to recall, the most parsimonious and simple approach is to adopt existing organisational systems. For this reason, laboratory settings can provide assessments of the organisational strategies that people use, as well as the extent to which certain strategies dominate under various conditions.

These results demonstrate people's dynamic ability to organise events. Not only can people organise events in multiple ways, they also cross-classify them in multiple organisations simultaneously, such that they can pivot from one organisation to another during retrieval. Much remains to be learned about the representational and processing mechanisms that make such flexible performance possible.

Manuscript received 28 August 1996
Manuscript accepted 30 August 1996

REFERENCES

Anderson, S.A., & Conway, M.A. (1993). Investigating the structure of autobiographical memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 1178–1196.

Ayres, J.S. [now Lancaster, J.S.] (1982). *Retrieval patterns in children's memory*. Masters thesis, Emory University.

Barsalou, L.W. (1988). The content and organization of autobiographical memories. In U. Neisser & E. Winograd (Eds.), *Real events remembered: Ecological and traditional approaches to the study of memory* (pp.193–243). Cambridge: Cambridge University Press.

Barsalou, L.W. (1995). Storage side effects: Studying processing to understand learning. In A. Ram & D. Leake (Eds.), *Goal-driven learning* (pp.407–419). Cambridge, MA: MIT Press/Bradford Books.

Barsalou, L.W., & Sewell, D.R. (1985). Contrasting the representation of scripts and categories. *Journal of Memory and Language*, *24*, 646–665.

Brown, N.R., Shevell, S.K., & Rips, L.J. (1986). Public memories and their personal context. In D.C. Rubin (Ed.), *Autobiographical memory* (pp.137–158). New York: Cambridge University Press.

Ceci, S.J., & Howe, M.J.A. (1978). Age related differences in free recall as a function of retrieval flexibility. *Journal of Experimental Child Psychology*, *26*, 432–442.

Chase, W.G., & Ericsson, A. (1981). Skilled memory. In J.R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp.141–190). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Conway, M.A. (1990a). Associations between autobiographical memories and contexts. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *16*, 799–812.

Conway, M.A. (1990b). *Autobiographical memory: An introduction.* Milton Keynes, UK: Open University Press.

Conway, M.A. (1990c). Conceptual representation of emotions: The role of autobiographical memories. In K.J. Gilhooly, M.T.G., Keane, & G. Erdos (Eds.), *Lines of thinking*, Vol. 2 (pp.133–143). New York: Wiley.

Conway, M.A., & Bekerian, D.A. (1987). Organization in autobiographical memory. *Memory & Cognition*, *15*, 119–132.

Conway, M.A., & Rubin, D.C. (1993). The structure of autobiographical memory. In A.F. Collins, S.E. Gathercole, M.A. Conway, & P.E. Morris (Eds.), *Theories of memory* (pp.103–137). Hove, UK: Lawrence Erlbaum Associates Ltd.

Crowder, R.G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Eldridge, M.A., Barnard, P.J., & Bekerian, D.A. (1994). Autobiographical memory and daily schemas at work, *Memory*, *2*, 51–74.

Graesser, A.C., & Mandler, G. (1975). Recognition memory for the meaning and surface structure of sentences. *Journal of Experimental Psychology: Human Learning and Memory*, *104*, 238–248.

Hasher, L., & Zacks, R.T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, *108*, 356–388.

Hyde, T.S., & Jenkins, J.J. (1969). Differential effects of incidental tasks on the organization of recall of a list of highly associated words. *Journal of Experimental Psychology*, *82*, 472–481.

Kolodner, J.L. (1978). *Memory organization for natural language database inquiry.* Research Report 142, Department of Computer Science, Yale University.

Kolodner, J.L. (1980). *Retrieval and organizational strategies in conceptual memory: A computer model.* Research Report No. 187, Department of Computer Science, Yale University.

Kolodner, J.L. (1983a). Maintaining organization in dynamic long-term memory. *Cognitive science*, *7*, 253–280.

Kolodner, J.L. (1983b). Reconstructive memory: A computer model. *Cognitive Science*, *7*, 281–328.

Kolodner, J.L. (1984). *Retrieval and organizational strategies in conceptual memory.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Lancaster, J.S. (1985). *Experimental investigations of the organization of memory for events.* Doctoral dissertation, Emory University.

Mandler, G. (1967). Organization and memory. In K.W. Spence & J.T. Spence (Eds.), *The psychology of learning and motivation, Vol. 1* (pp.327–372). New York: Academic Press.

Neely, J.H., Schmidt, S.R., & Roediger, H.L. III (1983). Inhibition from related primes in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 196–211.

Neisser, U. (Ed.). (1982). *Memory observed: Remembering in natural contexts.* San Francisco: W.H. Freeman & Co.

Neiser, U., & Winograd, E. (Eds.). (1988). *Remembering reconsidered: Ecological and traditional approaches to the study of memory.* New York: Cambridge University Press.

Nelson K. (Ed.). (1986). *Event knowledge: Structure and function in development.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Puff, C.R. (Ed.) (1979). *Memory organization and structure.* New York: Academic Press.

Reiser, B.J. (1983). *Contexts and indices in autobiographical memory.* Doctoral dissertation, Yale University, Cognitive Science Technical Report 24, Yale University.

Reiser, B.J. (1986). The encoding and retrieval of memories of real-world experiences. In J.A. Galambos, R.R., Abelson, & J.B. Black (Eds.), *Knowledge structures* (pp.71–99). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Reiser, B.J., Black, J.B., & Abelson, R.P. (1985). Knowledge structures in the organization and retrieval of autobiographical memories. *Cognitive Psychology*, *17*, 89–137.

Reiser, B.J., Black, J.B., & Kalamarides, P. (1987). Strategic memory search processes. In D.C. Rubin (Ed.), *Autobiographical memory* (pp.100–121). New York: Cambridge University Press.

Reitman, J.S. (1976). Skilled performance in Go: Deducing memory structures form inter-response times. *Cognitive Psychology*, *8*, 336–356.

Robinson, J.A. (1986). Autobiographical memory: A historical perspective. In D.C. Rubin (Ed.), *Autobiographical memory* (pp.100–121). New York: Cambridge University Press.

Roediger, H.L. III, Neely, J.M., & Blaxton, T.A. (1983). Inhibition from related primes in semantic memory retrieval: A reappraisal of Brown's (1979) paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 478–485.

Roenker, D.L., Thompson, C.P., & Brown, S. (1971). Comparison of measures for the estimation of clustering in free recall. *Psychology Bulletin*, *76*, 45–48.

Rubin, D.C. (Ed.). (1986). *Autobiographical memory.* New York: Cambridge University Press.

Salatas, H., & Flavell, J.H. (1976). Retrieval of recently learned information: Development of strategies and control skills. *Child Development*, *47*, 941–948.

Schank, R.C. (1982). *Dynamic memory: A theory of reminding and learning in computers and people.* New York: Cambridge University Press.

Schank, R.C., & Abelson, R.P. (1977). *Scripts, plans, goal, and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Schank, R.C., & Kolodner, J.L. (1979). *Retrieving information from an episodic memory.* Research Report 159, Department of Computer Science, Yale University.

Tulving, E. (1962). Subjective organization in free recall of ''unrelated'' words. *Psychological Review*, *69*, 344–354.

Tulving, E. (1964). Intratrial and intertrial retention: Notes toward a theory of free recall verbal learning. *Psychological Review*, *71*, 219–237.

Tulving, E. (1966). Subjective organization and the effects of repetition in multi-trial free recall verbal learning. *Journal of Verbal Learning and Verbal Behavior*, *5*, 193–197.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, *5*, 381–391.