Exp Brain Res (2002) 146:1–10
DOI 10.1007/s00221-002-1092-y

RESEARCH ARTICLE

Bettina D. Acuna · Jerome N. Sanes · John P. Donoghue

Cognitive mechanisms of transitive inference

Received: 18 October 2001 / Accepted: 11 February 2002 / Published online: 28 May 2002

Abstract We examined how the brain organizes interrelated facts during learning and how the facts are subsequently manipulated in a transitive inference (TI) paradigm (e.g., if A<B and B<C, then A<C). This task determined features such as learned facts and behavioral goals, but the learned facts could be organized in any of several ways. For example, if one learns a list by operating on paired items, the pairs may be stored individually as separate facts and reaction time (RT) should decrease with learning. Alternatively, the pairs may be stored as a single, unified list, which may yield a different RT pattern. We characterized RT patterns that occurred as participants learned, by trial and error, the predetermined order of 11 shapes. The task goal was to choose the shape occurring closer to the end of the list, and feedback about correctness was provided during this phase. RT increased even as its variance decreased during learning, suggesting that the learnt knowledge became progressively unified into a single representation, requiring more time to manipulate as participants acquired relational knowledge. After learning, non-adjacent (NA) list items were presented to examine how participants reasoned in a TI task. The task goal also required choosing from each presented pair the item occurring closer to the list end, but without feedback. Participants could solve the TI problems by applying formal logic to the previously learnt pairs of adjacent items; alternatively, they could manipulate a single, unified representation of the list. Shorter RT occurred for NA pairs having more intervening items, supporting the hypothesis that humans employ unified mental representations during TI. The response pattern does not support mental logic solutions of applying inference rules sequentially, which would predict longer RT with more intervening items. We conclude that the brain organizes information in such a way that reflects the relations among the items, even if the facts were learned in an arbitrary order, and that this representation is subsequently used to make inferences.

Keywords Mental model · Mental logic · Human psychophysics · Reaction time analysis · Symbolic distance effect

B.D. Acuna · J.N. Sanes · J.P. Donoghue (✉)
Department of Neuroscience, Box 1953, Brown Medical School, Providence, RI 02912, USA
e-mail: john_donoghue@brown.edu
Tel.: +1-401-8632701, Fax: +1-401-8631074

Present address:
B.D. Acuna, Section on Neurophysiology, NIMH, NIH, 49 Convent Dr., Rm B1EE17, MSC 4401, Bethesda, MD 20892-4401, USA

Introduction

It is not entirely known how the brain organizes information during learning, or whether this influences how facts are subsequently manipulated to infer new knowledge. For example, it would be advantageous if learned facts were stored in such a way that made the relationships among them explicit, since this, in turn, may facilitate problem solving and reasoning. We employed a transitive inference paradigm (TI; if A<B, B<C, then A<C) to examine knowledge manipulation from initial learning through the fully consolidated stage that allows abstract processes such as reasoning and problem solving to take place.

Previous studies have used TI paradigms to examine inferences after participants had learned the stimulus materials (Potts 1974). Since the main focus of those studies was the inference portion of the task, data such as reaction time (RT) during learning were not reported, and participants learned the stimulus materials through explicit, ordered rehearsal rather than elucidating the order via trial and error (Woocher et al. 1978; Henderson and Well 1985). Since behavioral patterns occurring during list learning can provide insight into how knowledge is organized and represented in the brain, our approach required participants to infer the ordering of the stimuli themselves. While the underlying structure of the stimuli was linear (A<B<C...<K), participants learned the order by exposure to arbitrarily chosen, individual pairs of adjacent stimuli (e.g., F vs E, B vs C) to impose as little

structure as possible. Specifically, we sought to determine whether participants memorize the individual pairs (e.g., F vs E, B vs C, etc.), or whether they integrate all the facts into one coherent representation (A<B<C...<K). A unified representation of the stimulus list would indicate that the brain organizes interrelated facts to make their relationships explicit, even if the facts are experienced only in isolation (in the current case, in pairs) and in an arbitrary order. We used RT to dissociate these two predictions. Previous work on paired associates has revealed that RT decreases during learning (Millward 1964; Kintsch 1965; Friedman and Clark 1967; Kakigi et al. 1985), suggesting that the learned pairs are stored individually. Alternatively, if premises become progressively unified into a single representation rather than stored individually, RT may increase as items are learned and incorporated into the representation. Indeed, recalling items from ordered lists requires more time for longer lists (Sternberg 1967; Juola et al. 1971). A similar pattern of increasing RT may be expected if participants add items into the list as they learn it, because the list would become longer until learning was complete.

In addition to examining how stimuli are mentally organized during learning, we characterized RT as participants solved TI reasoning problems after acquisition. We hypothesized that if the representation of the facts during learning reflected the inherent structure of the information, it would also influence subsequent manipulation of that knowledge. Prior data indicate that responses to TI problems are faster when more items from the learned list intervene between the current two test items (Moyer and Landauer 1967), termed the symbolic distance effect (SDE). However, the exact mental operations engaged during TI have not been identified yet. It has been suggested that reasoning may occur by either applying predetermined logical steps sequentially, as a schema (Braine 1998); by building and operating upon holistic mental representations (Johnson-Laird 1983); or by a combination of methods. In TI, the inferred knowledge may arise by repeatedly applying inference rules on the learned pairs. This process does not require a unified concept of the list, and it is akin to applying logical processes sequentially because it entails stepping through the learned stimulus pairs in order (e.g., A<B, B<C to determine A<C). Alternatively, reasoners may solve TI by manipulating a unified representation (A<B<C), which permits more flexible access to the data, such as at different points within the list, or perhaps at two locations in parallel to determine their relative position in the list. We observed TI on non-reinforced, novel combinations of the list items to determine whether an underlying structure created during learning influences subsequent mental processes, specifically inference. We examined whether TI occurs by serial search through individual stimulus pairs, which can be achieved simply by using the memorized stimulus pairs, or whether participants solve TI by manipulating a unified representation that takes into account the linear ordering. Portions of this work have been published in abstract form (Acuna et al. 1998).

Materials and methods

Participants

Twenty-five volunteers participated in this study; the first ten were included in the learning phase analysis. They were self-reported neurologically normal, between the ages of 18 and 25 years (6 women, 19 men, one left-handed, all others right-handed), and were recruited from the Brown University and local Providence communities. None presented obvious neurological deficits or other health-related problems. Each participant gave written informed consent according to established and approved guidelines by the Institutional Review Board for Human Investigations at Brown University, and each received modest monetary compensation.

Apparatus and procedures

Behavioral testing occurred in a dimly lit room. Participants sat about 1 m in front of a 19” color monitor that displayed the visual stimuli and a cursor representing hand position. In all tasks, the background color of the monitor was black. Tasks were run from a Macintosh Quadra 950 computer (Apple Computer Co., Cupertino, CA). To respond to visual stimuli, participants held a pen stylus with an embedded electromagnetic coil. The stylus was coupled to a Numonics Inc. digitizing tablet (Model 2020, Montgomeryville, PA), and it controlled a position feedback cursor on the monitor. The tablet had an active 60×60-cm measurement field, and stylus position was sampled at 100 Hz with an accuracy of 0.1 mm. Participants sat directly in front of the tablet as at a desk. They grasped the stylus with the dominant hand, and, by moving the stylus, participants controlled the position feedback cursor on the computer monitor.

Visual stimulus attributes

We used the same stimulus set of 11 items, represented here as A through K (Fig. 1A), for all behavioral tasks. The items consisted of various oval or rectangular shapes of different colors with areas ranging from 1.4 to 4.7 cm^2 (screen size) and subtending a mean visual angle of 1.5°. From A to K (Fig. 1A), the colors were red, yellow, purple, dark blue, lime green, brown-orange, fuchsia, light blue, white, orange, and green, respectively. Chromaticity was assessed with a Minolta Chroma Meter CS-100, and each stimulus attribute – width, height, luminance, x and y chromaticity values – was regressed on item location in the sequence. No stimulus attribute was described by a statistically significant linear fit ($P>0.05$; Fig. 2A–C). Thus, participants could not likely use those stimulus attributes to determine the location of an item in the sequence.

Trial events

Each trial was initiated by the appearance of a white annulus (“center hold-target”), subtending ~1.2° of visual angle, in the center of the computer monitor (Fig. 3A). After the participant had moved the position feedback cursor (~0.7° diameter) into the hold-target (Fig. 3B), the hold-target disappeared, and two items from the sequence appeared about 3.6° to the right and to the left of center, respectively (Fig. 3C). In addition, two white annuli (“movement targets”) appeared, each about 5.5° to the left or right of center. Participants moved the cursor towards one of the movement targets to indicate the “correct” item (Fig. 3D, see additional description below). Movement targets controlled for movement amplitude since they appeared at a fixed distance from the center, regardless of the size and shape of the test items, and they also controlled for speed-accuracy tradeoffs, because they were always the same size (1.8 cm diameter). The left-right position of test items varied randomly. Trials were aborted if either RT or move-

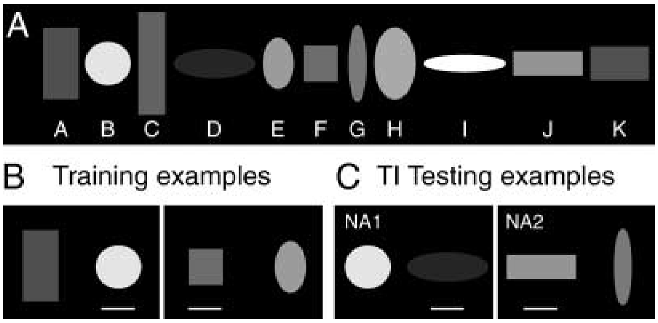

Fig. 1A–C Transitive inference sequence. **A** Eleven items were arranged in a fixed order as illustrated. Sequence items were presented only in pairs; the sequence as a whole is depicted here, together with letter referents (*A*–*K*), merely for illustrative purposes. For any two items, the "correct" choice was the item closest to *K*. Items were presented on a computer monitor with a black background. **B** Examples of two training pairs, one with an end-point (*left*) and one drawn from the sequence middle. The correct one of each pair (also in **C**) is underscored with a *white line*. **C** Examples of two test stimulus pairs, with non-adjacent (*NA*) stimuli (*left* NA1, *right* NA2)

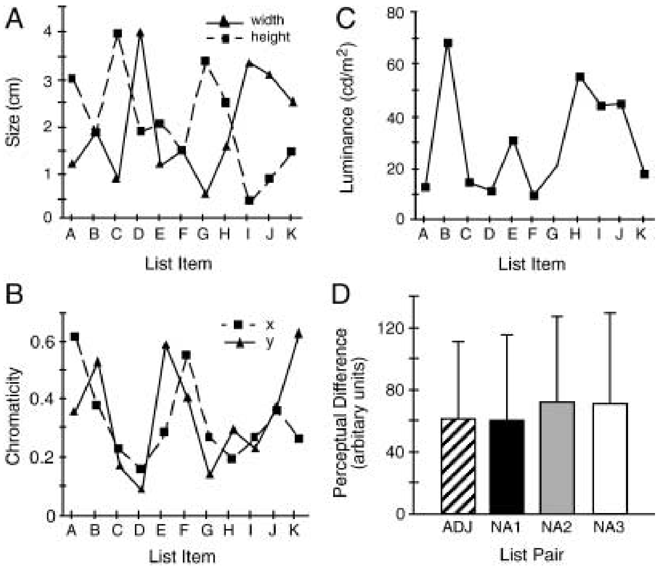

Fig. 2A–D Stimulus attributes. Abbreviations as in Fig. 1. **A** Height and width (cm) per item. **B** Chromaticity of *x* and *y* coordinates per item. **C** Luminance (candelas/m^2) per item. **D** Mean + SD perceptual difference (arbitrary units) between sequence items by trial type (Robertson 1977) [*thick cross hatch* adjacent pairs (ADJ), *shading* non-adjacent pairs with one to three intervening list items]. For additional details see "Materials and methods"

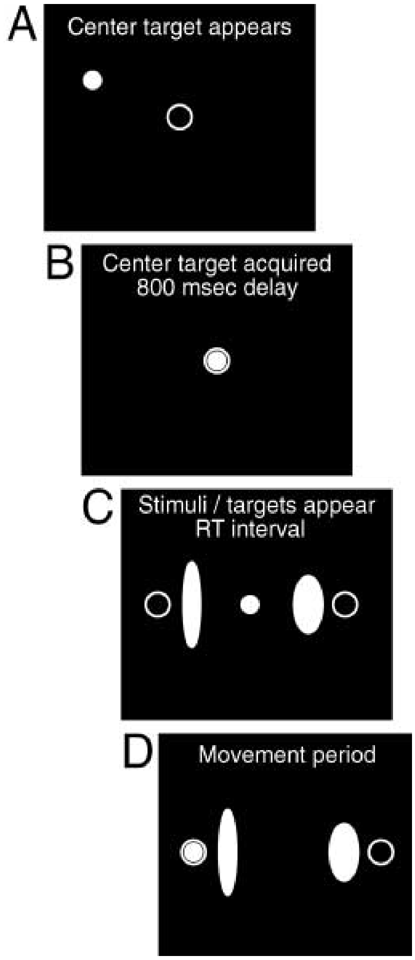

Fig. 3A–D Trial events. Typical appearance of the target and stimulus items on the video monitor during one trial. **A** At trial onset, a central hold-target appeared (*open circle*) and, to start a trial, participants would move the hand position cursor (*small filled white circle*) into the hold-target. **B** Alignment of the position cursor with the hold-target was required for 800 ms for the trial to proceed. **C** Two items from the sequence (*solid ovals*) appear along with their corresponding movement targets (*open circles*). Sequence items appeared in color (see Fig. 1), not in *white* as depicted here. **D** Participants chose the correct item and moved from the hold-target to the appropriate movement target as quickly as possible

ment time was longer than 2 s. The intertrial interval was 500 ms in all tasks, during which the video screen remained black.

Tasks

The experiment was divided into two phases: learning and testing, performed on the same day. Test stimuli were obtained from the same set of 11 visual shapes (Fig. 1A). The three tasks differed by the rule determining the correct choice between the paired stimulus items. For the "sequence" task, participants choose the item occurring "later" in the sequence, that is, nearest to item K. Since item A had the "earliest" position in the sequence, it could never be the correct choice, while K was correct whenever it appeared. Thus, given A vs B, B was correct; given B vs C, C was correct, and so on (Fig. 1B). The sequence task was performed during the learning and the testing phases. During the learning phase, participants learned the order of the list items by trial and error, receiving extensive practice on pairs of items that were adjacent in the list. During the testing phase, all participants performed the sequence task with the same stimuli, though under modified conditions (see below), as well as a height discrimination task. In the height discrimination task, participants applied the rule to "choose the taller item" to control for rule application during TI ("choose the later item"). For the height discrimination task, participants used the available visual information to choose the taller item regardless of position in the sequence or any attribute other than height. Only items that were adjacent in the sequence (e.g., C vs D) were presented together for height comparisons. The first ten participants also performed a "left-right alternation" task in the testing phase. This task served as a visual motor control, during which participants alternated moving the cursor to the left and the right movement targets, regardless of the items presented. This task used adjacent sequence elements.

Instructions to participants

The instructions were read to participants, and it was explained that they were required to learn a sequence, that items that were next to each other in the sequence would appear in pairs on the computer monitor, and that the correct answer in each trial was to choose the item occurring "later" in the sequence. A sample sequence of four novel black and white items that were unrelated to the experimental list was shown to illustrate the procedures. Participants were further told that they had to use trial and error to determine the correct answers, and that a beep would sound for every incorrect choice. Subsequently, the instructions detailed that the second task required choosing the taller object, while the third task (if given) simply required alternating between left and right items.

To elicit rapid and accurate RT, the instructions also indicated that participants should determine the correct stimulus as rapidly as possible and to move toward the appropriate movement target only after choosing an answer. This procedure served to minimize false starts and direction changing, each of which would have invalidated trials.

Sequence trial types

There were two types of trials that required knowledge of the sequence order: those having pairs of items that were adjacent (ADJ) in the sequence, such as B vs C, and stimulus pairs separated by one or more items, such as B vs D or D vs H, which we refer to as non-adjacent (NA). We used three types of NA trials – NA1, NA2, and NA3 – depending on whether stimulus pairs had one, two, or three intervening sequence items. With the 11-item sequence (A through K), we had 10 pairs of ADJ trials, seven NA1 pairs, six NA2 pairs, and five NA3 pairs. Each item pair yielded two instantiations, one for each left-right configuration on the display. There were no systematic differences in stimulus height, width, chromaticity, or luminance (Fig. 2A–C), and, furthermore, sequence pairs did not have systematic perceptual differences (Fig. 2D) as determined by a standard color-difference formula (Robertson 1977). Thus, RT differences among trial types cannot be explained by physical or perceptual stimulus attributes.

The reinforcement history of the items was equal except for the first and last items in the sequence (A and K). Since the sequence tasks required choosing the item occurring later in the sequence, choosing item A never resulted in positive reinforcement, while choosing item K always yielded positive reinforcement. In the sequence task, all other items were reinforced equally. Because of the unique reinforcement histories of A and K, they were not used when testing TI (NA pairs). Furthermore, feedback was provided only on ADJ pairs.

Learning phase

At the start of the learning phase, participants had no information regarding sequence order, but they learned it by receiving feedback on ADJ trials. An incorrect stimulus choice was signaled by a tone and, in this phase only, the trial was restarted until performed correctly. Participants saw only pairs of adjacent items during learning and were not shown the sequence in its entirety. Three criteria were used to verify that participants had explicit knowledge of the sequence. First, participants had to achieve ≥90% correct on ADJ trials. They were allowed to practice ADJ trials until they achieved ≥90% or until 4 h had elapsed. If they failed to reach 90% correct within 4 h, they were excluded from the study. Second, they had to correctly sketch the sequence from memory after reaching the 90% criterion. Finally, a test block of NA1 trials was given at the end of the learning phase to verify that participants were able to perform TI. On these trials, no feedback was provided, and participants could continue in the study if they performed at ≥70% correct on these trials. Acquisition required an average of about 850 (range 400–1,400) trials across participants. Three participants, not included in this report, failed to meet the criteria and their data were excluded from further analysis.

Testing phase

After fulfilling the learning criteria, participants performed two blocks of ADJ, three blocks of NA trials (one each for NA1, NA2 and NA3 stimulus pairs), two blocks of height comparison trials, and two blocks of left-right alternations (ten participants). ADJ blocks had 100 trials each. The NA1 block had 70 trials, the NA2 block had 60 trials, and the NA3 block had 50 trials, since each stimulus pair was repeated 10 times (five trials for each left-right configuration on the display). This totaled 460 trials for 15 participants that did not perform the left-right alternations and 540 trials for the 10 participants that performed the left-right alternations. Since the alternation task was viewed as a non-crucial control task, the final 15 participants did not perform it to reduce the total testing time. Each height discrimination and left-right alternation block contained 40 trials. For each participant, the task sets were ordered according to a Latin square design, and the trial order within each set was randomized. No reinforcement was provided on trials using NA pairs. The employed NA pairs never included the end-points A or K because these stimuli had a unique reinforcement history and therefore may have been solved by a trivial rule.

Data analysis

RT was measured from the time the sequence items appeared on the video screen to when movement velocity first exceeded 1 cm/s. Trials with errors, false starts, or mid-movement direction changes were excluded from further analysis. For the ADJ and height comparison tasks, the trial block that had the highest percent correct was analyzed, so as to compare peak performance across tasks.

ADJ pairs were divided into end-point ADJ and "conditional" ADJ pairs. End-point ADJ pairs consisted of the elements A with B and J with K and their left-right reversals on the display, while all the remaining ADJ pairs comprised the conditional ADJ group (BC/CB through IJ/JI). These stimulus pairs are termed conditional because the correct item depends on the relative sequence position of the other member of the pair, while end-point pairs can be considered a special case. For each participant, the median RT per pair was determined and used for subsequent analyses. RT during the learning phase was analyzed by dividing the first ten participants' learning data into early (50–66%), middle (67–83%), and late (84–100%) learning phases as determined by percent correct. To classify trials into early, middle and late learning phases, percent correct was determined for each block of trials during the learning phase. The chance level of performance, and therefore approximately the worst score a participant would get, was 50% because on each trial two choice stimuli appeared on the screen. The scores from 50% to 100% correct were divided a priori into three equal portions (50–66% correct, 67–83%, and 84–100%). Depending on the level of percent correct of each block, its trials were classified as early, middle, or late learning. Thus, the blocks of trials that resulted in 50–66% correct were classified as the early learning phase; those that resulted in 67–83% correct were classified as the middle learning phase; and those in which a participant achieved 84–100% correct comprised the late learning phase. We followed this procedure to take into account the participants' different numbers of trials and varying rates of learning. Finally, to determine the variability of RT across learning, standard deviations were calculated for every ADJ pair in each learning phase.

In general, the statistical analyses employed non-parametric tests (e.g., Kruskal-Wallis and Tukey-Kramer HSD) to compare

RT at different stages during learning, to compare RT pre- and post-learning, and to examine RT differences among tasks once the sequence order had been learned to criterion.

Results

Three main findings are reported here. First, as participants learned the order of items that were adjacent to each other in the list, RT increased while RT variability decreased. Second, during the testing phase, we observed shorter RT for NA pairs with more intervening sequence items than for NA pairs whose items were closer together in the sequence. Third, we found shorter RT for list pairs that were positioned closer to the sequence end-points.

Learning phase

Learning conditional ADJ pairs

To examine learning effects on conditional ADJ trials (ADJ pairs that did not include end-points A or K), the data were divided into early, middle, and late learning, according to a percent correct criterion. We found that RT increased significantly during learning (Kruskal-Wallis test, χ^2=34.25, $P\leq 0.0001$; Figs. 4, 5A). Specifically, middle and late learning trials elicited significantly longer RT than early learning trials (Tukey-Kramer HSD, q=2.58, $P\leq 0.05$, Figs. 4, 5A). RT for the conditional ADJ pairs did not differ significantly between middle and late learning, thereby indicating that RT reached a plateau after an initial increase. RT increases occurring during learning were accompanied by concurrent decreases in RT variability, as measured by the standard deviation of RT (RT_{SD}, Kruskal-Wallis test, χ^2=104.4, $P\leq 0.0001$, Fig. 5B). The RT_{SD} was higher during early and middle learning than during late learning (Tukey-Kramer HSD, q=2.58, $P\leq 0.05$).

RT obtained for conditional ADJ pairs during the early learning phase were significantly shorter than during testing (Tukey-Kramer HSD, q=2.58, $P\leq 0.05$; Fig. 5A), but RT during testing did not differ from that obtained during the middle and late learning phases. In contrast to the plateau in RT during later learning and the testing phases, the RT_{SD} exhibited a continuous decline throughout the learning phases that continued during testing. Further, the RT_{SD} was significantly higher during learning than during testing (Tukey-Kramer HSD, q=2.58, $P\leq 0.05$; Fig. 5B, compare open with closed symbols). This response behavior indicates that variability in performance continued to decline even after participants reached the learning criteria. In summary, RT to presentation of conditional ADJ pairs initially increased and then reached a plateau when participants fulfilled the learning criteria, and even while RT increased, the RT_{SD} decreased.

Learning end-point ADJ pairs

In contrast to conditional ADJ pairs, RT on trials with end-point ADJ pairs showed no significant changes during learning (Figs. 4, 5A). However, the RT_{SD} for end-point ADJ pairs decreased significantly (Kruskal-Wallis test, χ^2=36.9, $P\leq 0.0001$; Fig. 5B), specifically from early to middle and late learning (Tukey-Kramer HSD, q=2.63, $P\leq 0.05$; Fig. 5B).

RT for end-point ADJ pairs did not differ between the learning and testing phases (Fig. 5A). However, RT_{SD} declined during learning and testing; RT_{SD} during early and middle learning differed statistically from that obtained during the testing phase (Tukey-Kramer HSD, q=2.63, $P\leq 0.05$). Taken together, these results illustrate processing differences between end-point and conditional ADJ pairs, revealed primarily in the finding that RT increased for conditional ADJ pairs during learning, but not for end-point ADJ pairs.

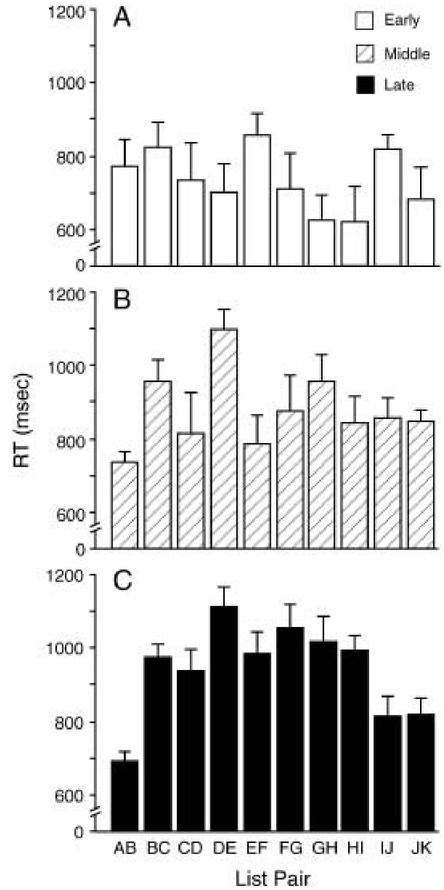

Fig. 4A–C Response times during sequence learning. Mean RT (+ SEM, N=10) for each stimulus pair during early **(A)**, middle **(B)**, and late **(C)** learning. Note that RT increases from early to middle to late learning, and that a response shape appears in late learning with the shortest RT occurring at the end-points and the longest RT in the middle of the sequence

Transitive inference testing

Errors during transitive inference testing

An analysis of errors committed during the testing phase for the sequence task provided additional evidence that

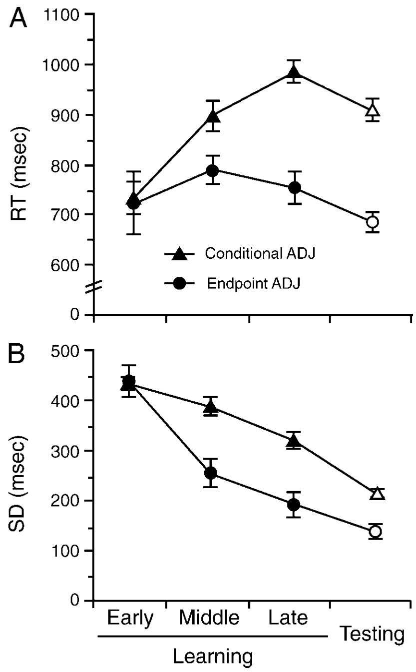

Fig. 5A, B Response times and variability for end-point and conditional adjacent pairs. Grouped RT and standard deviation of the RT (SD) during each learning phase (*closed symbols*) and during the testing phase (*open symbols*). **A** Means of ten participants' RT for conditional ADJ and end-point ADJ pairs (± SEM). **B** Means of the corresponding RT_{SD} for conditional ADJ and end-point ADJ pairs (± SEM)

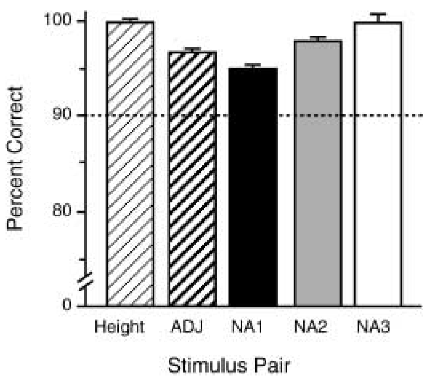

Fig. 6 Correct performance by trial type during testing. Histogram shows the mean percent correct (+ SEM, *N*=25) on the height comparison (*bar with thin hatched lines*) and sequence tasks (*remaining bars*). Sequence tasks consist of ADJ and NA pairs. Means for each task were above 90% correct, and participants received feedback only on ADJ trials

participants learned the sequence. Collectively, participants performed at ≥90% correct on all tasks (Fig. 6). Among sequence tasks, NA3 elicited the highest percent correct, while NA1 trials elicited the lowest (Fig. 6). Although the differences in error rates were small (Fig. 6), we observed a significant difference in percent correct across the type of NA trial (Kruskal-Wallis rank sum test, χ^2=17.31, *P*=0.05). Participants committed fewer errors for NA2 and NA3 trials than for NA1 trials

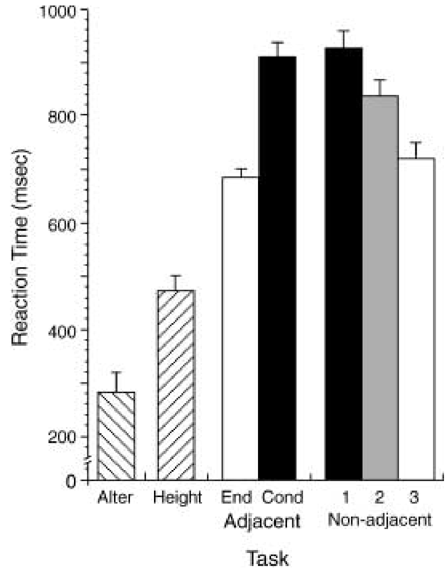

Fig. 7 Mean RT (± SEM) obtained for each task and trial type during testing (*Alter* alternation trials, *Height* height comparison trials, *End* end-point ADJ pairs, *Cond* conditional ADJ pairs). See "Results" for additional details

(Tukey-Kramer HSD test, *q*=2.62, *P*≤0.05). No other differences in percent correct among sequence trial types reached statistical significance.

Since percent correct and RT were both response variables, we assessed whether these measures covaried by correlating the percent correct performance with RT for NA pairs. We found a significant inverse rank order correlation (Spearman's *rho* = –0.24, *P*≤0.0005), indicating that, across NA pairs, RT increases accompanied decreases in percent correct.

Reaction time during the testing phase

Comparing non-adjacent, adjacent, and height pairs. Determining which of two stimuli is taller in the height comparison task required rule processing different from that required in the sequence tasks. To assess possible differences in cognitive mechanisms for these two rules, we compared RT among height, ADJ, and NA trials. The height task yielded significantly shorter RT than sequence trials (Fig. 7), even when compared to end-point ADJ trials, which elicited the shortest RT of all sequence trials (Tukey-Kramer HSD, *q*=2.89, *P*≤0.05). The alternation task, in turn, had significantly lower RT than the height comparison task (Wilcoxon rank-sum test, *Z*=–3.02, *P*≤0.005; Fig. 7). Thus, application of a simple rule, such as deciding the relative height of two visually presented objects, or implementing a simple movement strategy – alternation – yielded lower RT than TI or ADJ trials.

As noted, the end-points of the sequence had unique characteristics, since choosing item K always yielded positive reinforcement, whereas choosing item A never resulted in positive reinforcement. During learning, these attributes yielded lower RT for end-point than for conditional ADJ stimulus pairs. Similarly, during the testing phase, we found that RT for the end-point pairs AB and JK continued to elicit shorter RT than conditional ADJ, NA1, and NA2 stimulus pairs (Figs. 5A, 7; Tukey-Kramer HSD, q=2.89, P≤0.05).

Symbolic distance effect. The SDE predicts that trials with NA pairs with more intervening items will have shorter RT than trials with items closer together in the sequence (Moyer and Landauer 1967). The obtained RT data revealed the SDE insofar as the different types of NA pairs elicited systematically differing RT (Kruskal-Wallis test, χ^2=85.74, P≤0.0001; Fig. 7). RT scaled roughly linearly with number of intervening items, so that each additional intervening item decreased RT by an average of 86 ms (Fig. 7). Consistent with the SDE, RT for NA1 pairs was significantly longer than RT for NA3 pairs (Tukey-Kramer HSD, q=2.89, P≤0.05; Fig. 7), while RT for NA2 pairs was intermediate. The RT differences among NA pairs had between-subject reliability. A longer median RT for NA1 than for NA3 was found in 23 of the 25 participants (92%), while 20 of the 25 participants (80%) exhibited higher median RT for NA1 pairs than NA2 pairs and for NA2 than for NA3 stimulus pairs.

End-point proximity. Mental representations could index stimulus pairs relative to the sequence end-points as a mechanism to improve processing speed. To address this possibility, we assessed whether end-point proximity affected RT for NA and ADJ pairs. The analysis revealed that proximity to end-points significantly affected RT for ADJ (Kruskal-Wallis test, χ^2=70.77, P≤0.0001) and NA1 stimulus pairs (Kruskal-Wallis test, χ^2=42.94, P≤0.0001; Fig. 8). The stimulus pairs closer to the end-points, specifically AB, IJ, and JK, elicited the shortest RT of the ADJ pairs, while the middle pairs, DE through GH, elicited the longest RT (Fig. 8). The RT pattern for NA1 pairs also revealed an effect of end-point proximity. The two pairs closest to the reinforced end-point, GI and HJ, elicited significantly shorter RT than the remaining pairs, BD through FH (Tukey-Kramer HSD, q=2.98, P≤0.05; Fig. 8). For NA2 and NA3 pairs, however, there was no significant end-point effect, which may have been related to a number of possibilities such as having more intervening items than NA1 pairs or the relatively fewer number of pairs.

Proximity to the reinforced end-point. In addition to the finding that end-point proximity affected RT, we also examined whether proximity to the reinforced end-point K elicited shorter RT than proximity to the non-reinforced end-point A. ADJ pairs are not included in this analysis because participants always received performance feedback for these trials. Qualitatively, for each of the three NA trial types, the pair closest to the positively reinforced end-point K elicited shorter mean RT than the pair closest to the non-reinforced end-point A (e.g., the NA1 pair HJ had shorter RT than BD; Fig. 8). While these differences were not statistically significant for individual trial types, when we grouped NA pairs (comparing HJ, GJ, and FJ to BD, BE, and BF), we found that those closest to the positively reinforced end-point K yielded shorter RT than those closer to the end-point A (Wilcoxon rank-sum test, Z=–2.05, P≤0.05). However, this was a modest effect since individual trial types in each of NA1, NA2, and NA3 sets failed to show a significant effect of proximity to K.

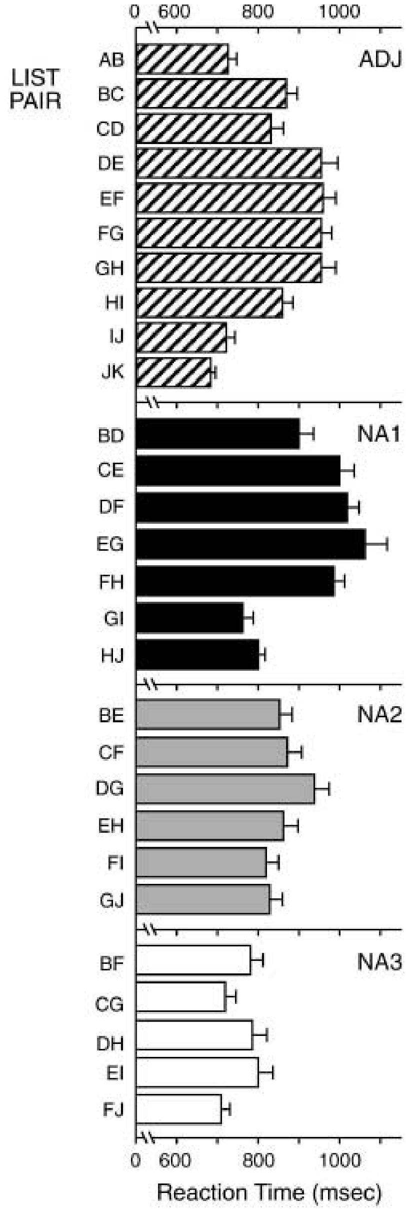

Fig. 8 Sequence pair RT during testing phase. Mean RT (+ SEM) are plotted for adjacent and NA1–3 pairs. Feedback was provided only for ADJ pairs. Abbreviations of pair types as in text

Interaction between symbolic distance and end-point effects. The results reported above suggest that the SDE and proximity to end-points both affected RT in a TI task. To determine whether these effects interacted, we performed a least squares multiple regression with nominal variables encoding the number of intervening items in each pair, the location of each pair in the sequence, and the participant to control for intersubject variability. ADJ, NA1, NA2, and NA3 stimulus pairs received scores of 0, 1, 2 and 3, respectively, depending on the number of intervening list items. Each pair's proximity to the end-points was scored by adding the number of items separating each member of the pair from the closest end-point. Thus, pairs in the middle of the sequence received the highest scores, while scores for pairs closer to the end-points A and K decreased linearly. For example, AB and JK received a score of 1 since A and K are the end-points (0 distance) and B and J are each one item away from an end-point. EF and FG had scores of 9, since E and G are four items from A and K, respectively, and F is separated by five items from both A and K. We found that the regression model significantly fit the data ($F_{(27,698)}$=24.94, $P \leq 0.0001$), and specifically that there were main effects of intervening elements in the NA sequence task ($F_{(1,698)}$=4.99, $P \leq 0.03$) and end-point proximity ($F_{(1,698)}$=187.53, $P \leq 0.0001$). In addition, we found an interaction between NA number and end-point effect ($F_{(1,698)}$=30.72, $P \leq 0.0001$).

Discussion

The experiment revealed three principal results: first, RT increased during learning even as variability decreased, contrary to paired associate learning (Millward 1964; Kintsch 1965; Friedman and Clark 1967). Second, during TI testing, item pairs with more intervening stimuli elicited shorter RT than closer sequence items. Third, pairs closer to the sequence end-points elicited shorter RT than those located centrally. The findings are consistent with the interpretation that participants build a unified mental representation from the individually learned premise pairs, and subsequently manipulate it to perform TI.

Organizing information in the brain

Previous studies have found decreases in RT as paired associate learning proceeds (Millward 1964; Kintsch 1965; Friedman and Clark 1967). Similar to these studies, we presented paired stimuli during learning. However, our stimuli could be represented as an ordered list as well, and participants had been instructed to determine the order. We tested whether the brain represents facts in such a way that reveals how the items are interrelated. The learned information could have been stored in several ways, such as encoding them as individual pairs or an integrated list, and TI problems could have been solved correctly either way. However, the stimuli consisted of a linear ordering (A<B<C...), so that a unified representation would be more efficient than representing each learned pair individually. In contrast to previous studies of non-ordered paired associates (Millward 1964; Kakigi et al. 1985), the current work showed that RT for ADJ pairs increased during learning. One interpretation of this result is that premises (item relationships) became incorporated into a unified and increasingly larger mental representation as learning progressed. Such a representation would likely require increasingly more time to process; indeed, previous work has found a positive correlation between RT and list length for item recall from ordered lists (Sternberg 1967; Juola et al. 1971; Corballis and Miller 1973). In these studies, participants explicitly memorized ordered lists, in contrast to the present work, but knowing the item position in the sequence was necessary for correct performance. Thus, the fact that list length affects RT indicates that lists of ordered items may be stored as a unit whether they are learned by explicit memorization or via trial and error, as is the case in the current work. This result suggests that the brain organizes information in a way that makes explicit the relationships among items, thereby allowing rapid inferences even between pairs that were never presented together during learning. The TI pairs, for example, were not presented during learning and never resulted in feedback of any kind, yet participants immediately solved these problems correctly.

An alternative interpretation for RT increases during learning is that proactive interference caused longer RT (Glass and Holyoak 1986). That is, stimulus similarity and high repetition rates interfered with recall of the correct answer and thereby increased RT as learning progressed. This explanation is unlikely since improvements in performance (percent correct) accompanied RT increases during learning; such improvements would not likely have occurred if answers were increasingly difficult to recall due to interference. Thus, the RT increase during learning suggests that it is more advantageous to store related pieces of information as a unified list than individually, despite the fact that participants were only exposed to individual pairs.

End-point effect

Our results demonstrate an end-point effect on RT during ADJ and NA1 trials, so that stimulus pairs closer to the end-points elicited shorter RT than middle pairs. Temporal primacy and recency effects (Rundus and Atkinson 1970) do not explain these results since the sequence pairs always appeared in a random order. However, the end-point pairs AB and JK have two properties that distinguish them from the other pairs: they have a unique reinforcement history and they also occupy unique positions in the sequence. The end-point pairs could, in theory, be solved in a different way than the conditional ADJ pairs are solved. For example, participants may recall that item K is always rewarded, and that item A is never correct. Using affective valence would not require recalling the sequence, and may thereby save time. However, the end-point pairs elicited significantly longer RT than the height comparison task, despite the fact that participants practiced end-point pairs extensively compared to the relative novelty of height comparison trials. This relatively long RT of end-point pairs suggests that participants recalled the sequence as a whole even for these pairs. The fact that AB and JK do have shorter RT than conditional ADJ pairs may indicate that participants more easily locate end-point items than items closer to the sequence middle, but that they nevertheless recalled the entire sequence while performing this task. Thus, recalling the sequence for end-point pairs may be a reason for the longer RT than height comparisons, while the uniqueness of the end-points in the sequence ordering, in turn, may cause these pairs to elicit shorter RT than conditional ADJ pairs.

RT decreased predictably the closer any given pair was to an end-point. Previous studies on TI have similarly concluded that items near end-points yield shorter RT (Trabasso and Riley 1975; Woocher et al. 1978). In addition, items near the end-point K may be searched earlier or more easily, since we found a weak effect of shorter RT for pairs closer to that end-point (Fig. 8). However, the decrease in RT near both end-points is not due to different reinforcement histories of items A and K, since the effect was seen in pairs that do not include end-points. For example, pairs BC and HJ have identical reinforcement histories as pairs in the middle of the sequence, yet BC and HJ tend to elicit shorter RT. The overall pattern of decreases in RT for ADJ and NA1 pairs near end-points suggests that participants may be searching items near end-points before searching items closer to the middle of the list. In addition, NA2 and NA3 pairs may not have shown an endpoint effect because the "spatial resolution" of NA2 and NA3 pairs may not have been sufficient to reveal an end-point effect in those trial types. Note that the number of items over which the end-point effect occurs (~3) approximates the number of intervening items in NA2 and NA3 pairs.

Methods of solving TI

Different solutions have been proposed regarding how reasoning problems are solved (Huttenlocher 1968; Johnson-Laird 1983; Adams 1984; Braine 1998). For the current TI problem, a mental logic approach (Braine 1998; O'Brien 1998) would predict that stimulus pairs having more intervening items should elicit longer RT if inference schemas are applied sequentially to intervening list items. However, we observed the opposite effect of longer RT with fewer intervening items, the SDE (Moyer and Landauer 1967). In fact, each additional intervening item saved about 86 ms in processing time. The SDE has also been found in comparisons of serial positions and geographic locations (Potts 1974; Maki 1981; Henderson and Well 1985). The current data fail to support the hypothesis that participants solve TI problems by serially processing intervening items or by sequentially applying mental logic rules (O'Brien 1998). Instead, the results suggest that accessing a sequence in more detail takes longer than comparing items that are further apart, such as by processing sequence items from both end-points inwards.

We found two main RT patterns when participants were asked to choose the item occurring closer to one end of a memorized list: the SDE and the end-point effect. Since item pairs with more intervening list elements appear to be discriminated more easily than those closer together, recalling additional detail in a mental representation appears to take time and does not seem to proceed serially from one end to another. A holistic representation of the list, rather than individual representations of learned pairs, would facilitate flexibly accessing learned information. Taken together, the results indicate that the brain organizes interrelated facts in a way that reveals not only the learned information, but also facts that were not specifically trained. Thus, one interpretation is that accessing and manipulating the unified structure that was created during learning may constitute the reasoning process in an inference task such as TI.

Acknowledgements National Institutes of Health grants AG10634 (J.N.S.), NS35376 (J.N.S.), and NS25074 (J.P.D.), and the James S. McDonnell Foundation (J.N.S.) supported this work. We would like to thank Carolyn K. Wang for help in conducting the experiments.

References

Acuna BD, Wang CK, Sanes JN, Donoghue JP (1998) Transitive inference mechanisms and the involvement of prefrontal cortex. Soc Neurosci Abstr 24:1178–1998

Adams MJ (1984) Aristotle's logic. In: Bower GH (ed) The psychology of learning and motivation: advances in research and theory. Academic Press, New York, pp 255–311

Attneave F, Curlee TE (1983) Locational representation in imagery: a moving spot task. J Exp Psychol Hum Percept Perform 9:20–30

Braine MDS (1998) Steps toward a mental-predicate logic. In: Braine MDS OD (ed) Mental logic. Lawrence Erlbaum, New Jersey, pp 273–332

Corballis MC, Miller A (1973) Scanning and decision processes in recognition memory. J Exp Psychol 98:379–386

Friedman MP, Clark MC (1967) Paired-associate learning in a simple adaptive environment. Psychon Sci 7:409–410

Glass AL, Holyoak KJ (1986) Cognition. Random House, New York

Henderson JM, Well AD (1985) Symbolic comparisons with and without perceptual referents: is interval information used? Mem Cognit 13:176–182

Howard RW (1980) Category use in abstract mental comparisons. Q J Exp Psychol 32:625–633

Huttenlocher J (1968) Constructing spatial images. Psychol Rev 75:550–560

Johnson-Laird PN (1983) Mental models. Harvard University Press, Cambridge, MA

Juola JF, Fischler I, Wood CT, Atkinson RC (1971) Recognition time for information stored in long-term memory. Percept Psychophys 10:8–14

Kakigi S, Shiwa S, Matsuda T, Mori T (1985) Changes in contingent negative variation (CNV) as a function of paired associate learning. J Psy Res 27:45–49

Kintsch W (1965) Habituation of the GSR component of the orienting reflex during paired-associate learning before and after learning has taken place. J Math Psychol 2:330–341
Kosslyn SM, Murphy GL, Bemesderfer ME, Feinstein KJ (1977) Category and continuum in mental comparisons. J Exp Psychol Gen 106:341–375
Maki RH (1981) Categorization and distance effects with spatial linear orders. J Exp Psychol Hum Learn Mem 7:15–32
Millward R (1964) Latency in a modified paired-associate learning experiment. J Verb Learn Verb Behav 3:309–316
Moyer RS, Landauer TK (1967) Time required for judgments of numerical inequality. Nature 215:1519–1520
O'Brien DP (1998) Mental logic and irrationality: we can put a man on the moon, so why can't we solve those logical reasoning problems? In: Braine MDS OD (ed) Mental logic. Lawrence Erlbaum, New Jersey, pp 23–43
Pohl RF (1990) Position effects in chunked linear orders. Psychol Res 52:68–75
Potts GR (1974) Storing and retrieving information about ordered relationships. J Exp Psychol 103:431–439
Robertson A (1977) The CIE 1976 color-difference formulae. Color Res Applic 2:7–11
Rundus DL, Atkinson RC (1970) Rehearsal processes in free recall: a procedure for direct observation. J Verb Learn Verb Behav 9:99–105
Seamon JG (1972) Serial position effects in probe recall: effect of rehearsal on reaction time. J Exp Psychol 96:460–462
Sternberg S (1967) Retrieval of contextual information from memory. Psychon Sci 8:55–56
Trabasso T, Riley CA (1975) On the construction and use of representations involving linear order. In: Solso RL (ed) Information processing and cognition: the Loyola symposium. Lawrence Erlbaum, Hillsdale, NJ, pp 381–410
Woocher FD, Glass AL, Holyoak KJ (1978) Positional discriminability in linear orderings. Mem Cognit 6:165–173