Citation: Murphy, R. A., Mondragon, E. & Murphy, V. A. (2008). Rule learning by rats. Science, 319(5871), pp. 1849-1851. doi: 10.1126/science.1151564

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/22060/

Link to published version: https://doi.org/10.1126/science.1151564

Rule Learning by Rats

Summary

Rats can transfer learning of a three-element rule-based sequence to novel exemplars.

Authors

Robin A. Murphy[1*], Esther Mondragón[1*], Victoria A. Murphy[2]

[1]University College London, [2]University of Oxford, United Kingdom

*To Whom Correspondence should be addressed email: robin.murphy@ucl.ac.uk, e.mondragon@ucl.ac.uk

Using rules extracted from experience to solve problems in novel situations involves cognitions such as analogical reasoning and language learning and is considered a keystone of humans' unique abilities. Non-primates, it has been argued, lack such rule transfer. We report that *rattus norvegicus* can learn simple rules and apply them to novel exemplars. Rats learned that sequences of stimuli consistent with a rule (e.g., XYX), were different from other sequences (e.g., XXY, YXX). When novel stimuli were used to construct sequences that did, or did not, obey the previously learned rule, rats transferred their learning. Therefore, rats, like humans, can transfer structural knowledge from sequential experiences.

The ability to extract generalizable rules from specific experiences is a fundamental attribute of human higher cognitive functioning (1, 2). For instance, human language learning relies on learning grammatical rules that allow the English speaker to discriminate; ‘The dog bit the woman’ (Subject-Verb-Object) vs ‘Bit the woman the dog’ (Verb-Object-Subject). Rules importantly can also be transferred to newly encountered items contributing to the understanding and production of new sentences. Cross-linguistic differences in grammatical rules serve to highlight the challenge facing the developing child, that phrase structure rules must be learned. Language is only one of many cognitive domains where problems of this type are present raising the issue of the nature of the underlying cognitive substrates supporting rule learning in general.

Rule learning has been investigated in different species and stages of human development (3, 4). Pre-linguistic infants (5), primates (6) and even some birds (7, 8) can learn rule-like temporal structures. Starlings, for example, are able to learn sequential structures made up of segments of bird song and can recognize whether patterns are consistent with a training set. Although the question of whether the birds used recursion to solve the discrimination is controversial, they do appear to use a form of rule (8,9).

In this paper we studied the ability of rats to learn and transfer rules. We developed a procedure to test rule-transfer learning in the rat, an animal it has been argued might have a limited or perhaps even no such parallel capability to that used by humans (3, 5). Rats were trained with three-element sequences (e.g., ABA, BAB) paired with food. We asked whether they would learn something about the overall pattern of cues (i.e., XYX). There are at least three cognitive abilities required in order to learn this rule and apply it in novel situations. The first is an ability to learn that a sequence of cues signals food (10, 11). Second, animals need to discriminate sequences paired with food from those that are not. Rats can learn two-element sequences, A followed by B rather than B followed by A (12), as signals for food. The third is the ability to transfer to novel instances to test whether the rats learn a generalized rule rather than simply the trained instances. The transfer involved novel sequences that were either consistent or inconsistent with the previously trained rule, but that used stimuli that had never been exposed or trained.

Discrimination between stimuli is not a fixed ability as animals usually show some generalization to novel items. For example, to the extent that a new stimulus C has more perceptual similarity to A than to B, animals will treat C like it is A (10, 13). Furthermore, similarity is not restricted to perceptual features. Animals will treat two different cues with the same consequences (e.g., paired with food) as similar (14). Finally, animals will sometimes treat stimuli that have the same relative position on a stimulus dimension (e.g., greater than, brighter than) as similar (15). In our experiment the stimulus changes within and between sequences were not monotonically distributed along a dimension. There were no stimulus specific perceptual cues or common consequences that previous research might have predicted could have allowed generalization of behaviour.

The rules were based on experiments that have shown that prelinguistic human infants behave as if they learn rules embedded in sequences created from auditory phonemic cues (5, 16). We used patterns of visual (Experiment 1) and auditory (Experiment 2) stimuli that obeyed a similar rule.

The cues in Experiment 1 consisted of three-element sequences comprising short 10-second exposures to two visual cues: a bright light (A), or darkness (B) presented in a dimly lit environment. The rats were divided into three groups with each group receiving Pavlovian conditioning for food following sequences that obeyed one of three rules. One group received food with the XYX rule (i.e., ABA and BAB), another group with the XXY rule (i.e., AAB and BBA) and finally a third group with the YXX rule (i.e., BAA and ABB). Each group also received trials with the other two sequences but without food. Would each group learn their rule for food and discriminate the reinforced from non-reinforced sequences? We compared the rats' anticipatory responding for food during the third element of each sequence but before the food was available. The animals were able to discriminate the sequences by the end of the experiment (Table 1). They came to respond more at the end of sequences that signalled food than to trials of either of the other two rules and there was no evidence for differential responding to these two non-reinforced rules (17). Since A and B immediately preceded both food delivery and the absence of food, animals could not use the identity of the final stimulus to solve the discrimination. One possible solution involves the rats using unique pairs of stimuli that might have differentially signalled reinforced from nonreinforced sequences. Rats trained with XXY as the reinforced sequence might only have used the first pair of stimuli as a unique cue, while those trained with YXX might have used the final pair of stimuli; finally, the group whose reinforced sequence was XYX could have learned to use a pair formed by the first and third stimuli. However, if animals did make use of different stimulus combinations in each group we would expect discrimination levels to differ because each of these discriminations would have involved quite different levels of difficulty. For instance, using a strategy based on the first two stimuli to discriminate XXY from the other patterns would imply that the third stimulus is uninformative. The third stimulus would effectively delay reinforcement, and weaken the discrimination, relative to reinforcement following YXX in which food was contiguous with the unique stimulus pair. Overall, the notion that only stimulus pairs were crucial for learning does not account for the similar learning observed in all three groups. Discrimination performance did not differ among groups (17), suggesting that animals were anticipating food delivery by learning the whole triplet stimulus pattern. Rats had to use all three stimuli and their position in the sequence, a form of rule learning, to solve the discrimination.

These results demonstrate that rats can learn three-element sequences, but it is possible that one method of solving the task involved memorising the two sequences paired with food. Perhaps multiple instances of the rule (e.g., ABA and BAB) were treated as behaviourally similar but not perceived as examples from a broader category.

In a second experiment we used a transfer test to explore this possibility. We tested whether rats could generalise learning of a rule to novel stimuli that were either consistent or inconsistent with the trained rule. We used a similar procedure to Experiment 1, but because there were no differences in learning the three rules in Experiment 1 we trained all subjects with rule 1 (XYX) using two auditory pure tones (e.g., A = 3.2 kHz and B = 9 kHz). Auditory cues were used because they allow a wider range of cues for the transfer test. Animals received food at the end of ABA and BAB sequences, but not following BBA, AAB, BAA or ABB. After acquisition, we presented them with transfer stimuli comprised of two novel pure tones (e.g., C = 12.5kHz and D = 17.5kHz). The stimuli were counterbalanced so that the stimuli in the roles of A, B and C, D were reversed for half the animals and were chosen to ensure that no common frequency relation was present between the pairs. If rats had simply learned something specific about the reinforced elements ABA they should have been unable to choose CDC and DCD over CCD, DDC, CDD and DCC. The time the rats kept their head in the food trough during the final element of the sequence was used as a measure of learning. The results of the transfer test are presented in Fig. 1 excluding two rats that failed to learn the initial discrimination. More anticipatory behaviour for food was exhibited during sequences that were consistent with the previously learned rule even though the rats had never been presented with these particular instances and there was no food presented during the test. As with Experiment 1, animals responded more to the rule-consistent sequences compared with both inconsistent sequences. There were no differences in responding to the two inconsistent rules (17). To test that the rats remembered the trained sequences we presented the original stimuli again, without any food reinforcement or any retraining and found that they could still apply the rule to the original stimuli and consequently respond more to the previously reinforced sequences (See Fig. 1).

The design of these experiments rules out a number of simple explanations of the rats' behavior. They could not have solved this discrimination simply by learning the final element of the sequences, because these elements were the same for both reinforced and non-reinforced sequences. Similarly, they could not use pairs of stimuli embedded in the sequences to solve the discrimination as this would lead to differing levels of performance in Experiment 1. The result of Experiment 1 could be explained by the animals memorising the sequential configuration of the reinforced triplets. Experiment 2 showed that the rats transferred what they had learned about the sequences to novel instances. This transfer could not be achieved by memorizing the specific patterns. The counterbalancing of the exemplars of each rule prevented learning of a particular pitch or timbre pattern instead of the sequence pattern. Animals were trained with both High - Low - High and Low - High - Low frequency changes and therefore no simple pitch change can account for the original or transfer results. Pitch changes along the frequency dimension were not fixed, did not follow a continuum and the transfer cues were outside of the training range. Transfer had to be based on the entire relational sequence, resulting in rule-based behavior.

Finally, it remains possible that the animals learned the relational sequences as two separate but unrelated relational patterns and not as instances of the same category. However, during transfer rats treated both patterns in the same way and equally different from all other instances belonging to a different category (17). Hence, an account that integrates both patterns as instances of the same category is parsimonious and in line with data suggesting rule governed performance.

These experiments show that rats can learn simplified rules and apply these rules to novel stimuli. Even if the cognitive mechanism that allows the rats' behaviour is found to be dissimilar to that found in humans, it does seem that both species have evolved means of solving similar information problems, in this case transfer of overall sequence relations (see also 5, 6, 9). These experiments suggest that some rule governed behaviour is present in rats and well-understood principles of conditioning provide a perspective on our most human of cognitive abilities.

References and Notes

1. S. Pinker, *Science* **253**, 530 (1991).

2. D.C. Penn, D.J. Povilenni, *Ann. Rev. Psych.* **58**, 97 (2007).

3.C.M. Conway, M.H. Christiansen, *Trends in Cog. Sci.* **5**, 539 (2001).

4. J.L. Elman, et al. *Rethinking innateness: A connectionist perspective on development.* (MIT Press: Cambridge, MA 1996).

5. G.F. Marcus, S. Vijayan, S. Bandi Rao, P.M. Vishton, *Science* **283**, 77 (1999).

6. M.D. Hauser, D. Weiss, G. Marcus, *Cognition* **86**, 15 (2002).

7. W.T. Hebranson, C.P. Shimp, *Learn. & Behav.* **31**, 98 (2003).

8. T.Q. Gentner, K.M. Fenn, D. Margoliash, H.C. Nusbaum, *Nature* **440**, 1204 (2006).

9. M.C. Corballis, *Cog. Sci.* **31**, 697 (2007).

10. I.P. Pavlov, *Conditioned reflexes* (Dover Press: London 1927).

11. N.J. Mackintosh, *The psychology of animal learning* (Academic Press: London 1974).

12. R.A. Murphy, E. Mondragón, V.A. Murphy, N. Fouquet, *Behav. Proc.* **67**, 303 (2004).

13. N. Guttman, H.I. Kalish, *J. Exp. Psych.* **51**, 79 (1956).

14. R.C. Honey, A. Watt, *J. Exp. Psych. Anim. Behav. Proc.* **24**, 325 (1998).

15. K.W. Spence, *Psychol. Rev.* **44**, 430(1937).

16. J. Saffran, E. Aslin, E. Newport. *Science* **274**, 1926 (1996).

17. Materials, methods and procedural details as well as additional statistical analyses are available as supporting material on *Science* online.

18. We would like to thank E. Alonso, D.C. Penn, J.Tennenbaum for useful discussions and A.G.Baker and D.Shanks for comments on the manuscript. This research was supported by a grant from the BBSRC (UK).

Table 1. Illustration of the three rules and mean rate (per min.) of food tray entries in each group to both reinforced (RF) and non-reinforced (NRF) rules on the first and last block of training. Each group received food following one of the three rules (XYX, XXY or YXX) and no food following the other two rules. While no differences were found between the three groups or between reinforced and non-reinforced rules during the first two blocks of training trials [Fs (2,12) < 1], by the last block of training trials the main effect of reinforced *vs.* non-reinforced rule was statistically significant in all three groups [F(1,12) = 8.21, p < .01], although neither the main effect for the different groups or the interaction were statistically significant [Fs(2,12) < 1] suggesting that the effect was similar with all three rules.

Trained Rule		Group 1 Rule 1 (XYX) ABA, BAB		Group 2 Rule 2 (XXY) BBA, AAB		Group 3 Rule 3 (YXX) BAA, ABB	
		RF	NRF	RF	NRF	RF	NRF
Block of Learning Trials	First	12.66	12.58	9.14	9.81	13.94	13.97
	Last	25.97	24.18	24.47	22.97	32.08	29.06

Fig. 1. Mean time (centiseconds) with head in the food trough during the final 10 seconds of a sequence subtracting the 10 second prestimulus interval for that trial. Response rates are shown separately for consistent sequences and inconsistent sequences and the analysis confirmed that there was more responding to consistent than inconsistent sequences during the transfer test in the absence of food [$F(1,13) = 5.35$, $p < .05$] and during a test of the original stimuli [$F(1,13) = 4.87$, $p < .05$].

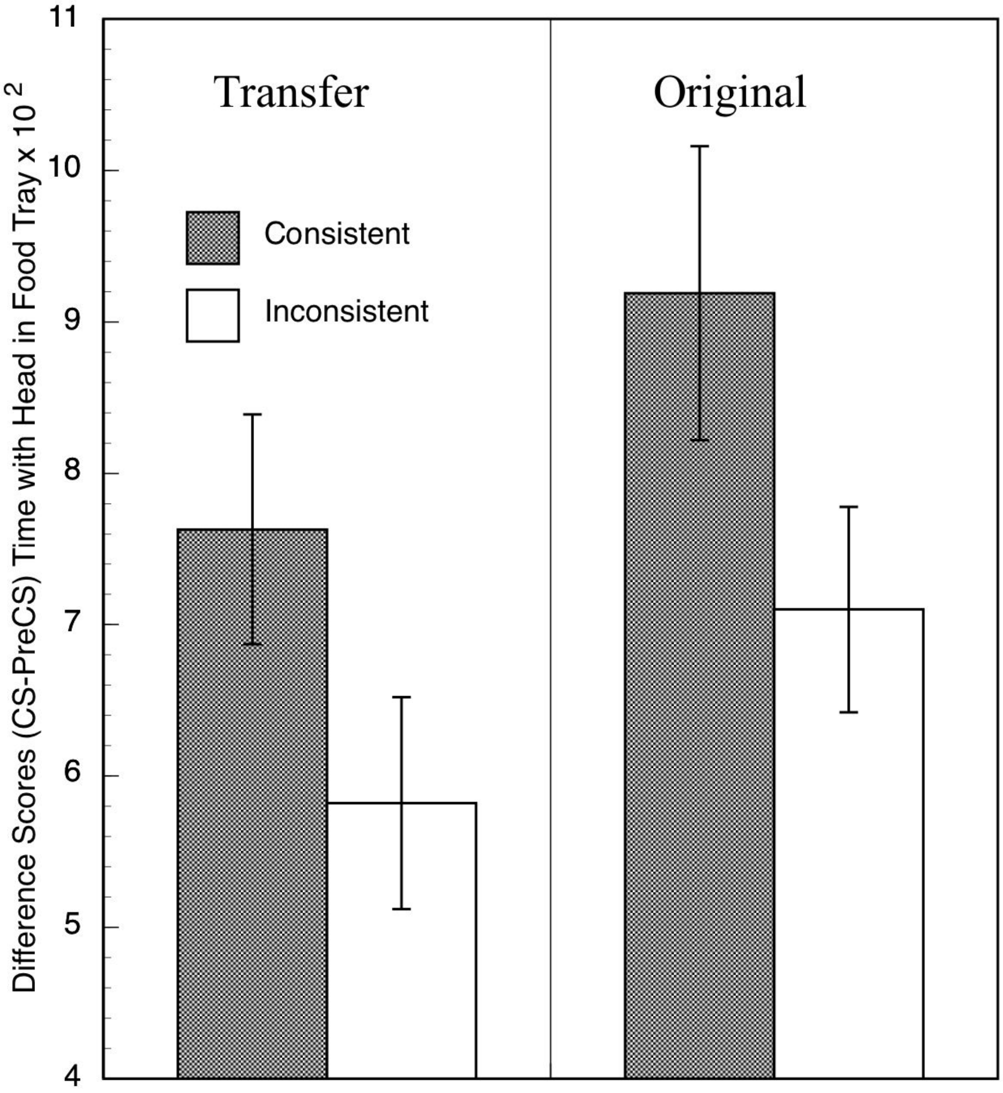
Transfer
Original
Consistent
Inconsistent
Difference Scores (CS-PreCS) Time with Head in Food Tray x 10 2
11
10
9
8
7
6
5
4

Supporting Online Material

Experiment 1

Method

Subjects. Fifteen male Lister hooded rats (Charles-River, London) initially weighing 302-352 grams were food restricted to 85% of their free-feeding weights for the duration of the experiment. Housing was maintained on a 12 hour light dark cycle with training taking place during the light part of the cycle.

Apparatus. Eight identical conditioning chambers with internal dimensions 30.5 cm X 21 cm X 24.1 cm were used. The chambers were housed in sound and light attenuating shells. The interiors consisted of a floor of metal bars 4.8 mm in diameter and 6.4 mm apart. The ceiling and two walls were metal with the two other walls in clear polycarbonate. Two visual cues were used in this experiment: darkness involved turning off the houselight (20V) and light which involved illuminating a bright stimulus light (28V).

*Procedure.*Training consisted of a single food retrieval pre-training phase in which food pellets were delivered on a variable interval 60 second schedule followed by training with the stimulus sequences. Each training session consisted of 60 trials presented with a variable intertrial interval (VI) of 60 seconds. Training continued for 24 sessions. Each trial consisted of a sequence of three 10 second presentations of one of two stimuli separated by a 0.5 seconds interval. The stimuli consisted of two visual cues, the illumination of a stimulus light or darkness (houselight off). All animals were presented with 10 trials of each of 6 different sequences, ABA, BAB, BBA, AAB, BAA and ABB, that is, 20 trials of each rule

(XYX, XXY, and YXX). Three groups of rats (n = 5) differed with respect to which of the three rules ended with the delivery of a single food pellet. Animals in Group XYX received food after ABA and BAB trials; AAB and BBA trials were followed by food in Group XXY; Group YXX received food after BAA and ABB trials.

The number of head entries into the food trough was recorded when a photo-electric beam crossing the entrance of the trough was broken. Responses were recorded during the final stimulus of each three-element sequence (the CS interval) and during the 10 seconds preceding the first element (the preCS interval).

Results

The results of the training were analysed by averaging difference scores (CS interval - preCS interval) in 6 four-session blocks. A mixed ANOVA comparing responding across the blocks supports the conclusion that differences in responding to the reinforced and nonreinforced sequences emerged over the course of training. The only statistically significant results from this analysis were the main effect of training blocks [F(5,60) = 25.84, p < .001)] and the interaction between blocks and reinforced vs nonreinforced rule [F(10,120) = 1.97, p < .05)]. The interaction reflects the fact that there was no differential responding at the start of training [F(1,12) < 1] whereas during the final two blocks of training there was a significant difference in responding in all groups, the rats responded more on the trials that were to be followed by food than on those trials that did not [F(1,12) = 8.21, p < .01]. None of the effects involving rule type (XYX, XXY or YXX), rule by trial type [F(2,12) < 1] or rule

(F(2,12) < 1] was significant suggesting that the ability to discriminate the food-paired sequence from the nonfood-paired sequence was similar regardless of which sequence was paired with food. It is possible that rats were solving this discrimination by withholding responding to only one of the two types of non-reinforced sequences. To test this possibility we subjected these response rates to further statistical scrutiny by conducting a separate analysis directly comparing responding to the reinforced sequence to each of the non-reinforced sequences (means = 27.50, 25.87, and 24.94 for the consistent and the two inconsistent sequences respectively). Animals responded more to the reinforced than to either nonreinforced sequence [$F(1,14) = 4.84$, $p < .05$ and $F(1,14) = 8.13$, $p < .05$] and there were no differences in responding between the pairs of non-reinforced sequences [$F(1,14) = 1.12$, $p > .30$]. Therefore the discrimination does not seem to reflect some type of strategy that encouraged differential responding to the inconsistent sequences.

Experiment 2

Method

Subjects. Sixteen male Lister hooded rats (Charles-River, London) initially weighing 310-330 grams were food restricted to 85% of free-feeding weights. To accommodate developmental growth 5% was added to the goal weight every 3rd week following the start of the experiment. Housing was maintained on a 12 hour light dark cycle with training taking place during the light part of the cycle.

Apparatus. The apparatus and stimuli used in Experiment 2 were the same as in Experiment 1 except that the visual stimuli were replaced with pure tones generated using a tone generator and presented through a speaker mounted on the ceiling of the chamber. Four different tones were generated [A = 3.2KHz; B = 9KHz; C = 12.5KHz; D = 17.5KHz tones].

Procedure.

Training. In an initial session the animals were trained to find the food using the procedure described in Experiment 1. Training on the sequences continued for 75 sessions. Each one hour session involved 18 trials (3 of each type of trial, 6 consistent and 12 inconsistent). Trials appeared on a VI 155 second schedule. Each trial comprised a 10 seconds presentation of a CS1, followed by a 10 seconds CS2 and finally by a 10 seconds CS3 presentation. Each stimulus was separated from the following one by a gap of 2 seconds. The termination of CS3 was followed immediately by food on consistent sequences. The reward consisted of 4 food pellets. Food pellets followed ABA and BAB. For half the rats A and B were a 3.2 and a 9 kHz tone respectively while for the other half of the animals A and B were 12.5 and 17.5 kHz tones. The amount of time the animals kept their head in the food trough, a reliable measure of conditioned behaviour to auditory cues, was recorded during the final stimulus of each compound and during a 10 second PreCS period. Two rats were removed from the analysis of the test data for failing to learn this initial discrimination (> 3 Standard Deviation away from the mean).

Test. On the first session following training animals were exposed to the same consistent and inconsistent sequences except that the stimuli were changed. The animals trained with the 3.2 and 9kHZ tones now received non-reinforced exposure to stimulus sequences of 12.5 and 17.5 kHz. The other half of the animals were exposed to the other two tones. The procedure and number of stimuli (18 trials in total) presented was identical to training except for the change in the nature of the two cues and the fact that no food was delivered. On the day after this transfer test a second test of the original training stimuli was conducted. The original training stimuli were presented using the same procedure as during training but, as with the transfer test, no food pellets were delivered at the end of the consistent rule sequences.

Results.

Animals responded more to the third element of consistent than to inconsistent sequences. Responding was also higher to the original training stimuli than to the novel test stimuli. Fig. 1 shows mean difference response time rates to the consistent and inconsistent stimuli during the transfer and original stimulus tests. An overall analysis of variance (ANOVA) conducted on these data with test (Transfer, Original) and stimulus (Consistent, Inconsistent) as variables, found a statistically significant main effect of Consistent vs Inconsistent sequences [$F(1,12) = 12.81$, $p < .01$], and a marginal effect of test [$F(1,12) = 3.96$, $p < .070$]. Although no other effects were reliable, including the interaction between test and stimulus, two separate direct tests of learning during the tests were conducted. These analyses found separate reliable discriminations between consistent and inconsistent sequences for both tests [$F(1,13) = 5.35$, $p < .05$, $F(1,13) = 4.87$, $p < .05$ for the transfer and original stimulus test, respectively]. As with Experiment 1 we also looked at the data more closely to ensure that responding to the consistent sequences was greater than responding to both types of inconsistent sequences. The mean difference response time rates during the transfer test were 775.93, 564.69, and 529.46 centiseconds for the consistent and the two types of inconsistent sequence respectively; during the test with the original stimuli means were 903.36, 696.96, and 713.14 centiseconds for consistent and the two types of inconsistent sequence respectively. An overall analysis of variance of these data found a main effect of test type [F(1,13) = 5.23, p < .05] indicating that responding was higher during the test of the original cues, and a significant effect of sequence type [F(1,13) = 8.77, p < .001] with no other interactions of main effects, suggesting differential responding to the sequences was similar during both tests. We conducted a further analysis showing that responding to consistent sequences was greater than to either inconsistent sequence [F(1,13) = 10.47, p < .01 and F(1,13) = 29.16, p < .001 respectively]. Finally, an analysis of the crucial transfer test data showed that rats responded more to the consistent sequences than to either type of inconsistent sequence [F(1,13) = 4.51, p < .05 and F(1,13) = 6.14, p < .01]. Furthermore, there was no evidence of any difference in responding to the two types of inconsistent sequence [F(1,13) < 1].