Current psychology letters

Behaviour, brain & cognition

17, Vol. 3, 2005 | 2005

Varia

Implicit learning in a prediction task: Neither abstract nor based on exemplars

Arnaud Destrebecqz and Axel Cleeremans

Electronic version

URL: <http://journals.openedition.org/cpl/470>

DOI: 10.4000/cpl.470

ISSN: 1379-6100

Publisher

Centre PsyCLÉ

Electronic reference

Arnaud Destrebecqz and Axel Cleeremans, « Implicit learning in a prediction task: Neither abstract nor based on exemplars », *Current psychology letters* [Online], 17, Vol. 3, 2005 | 2005, Online since 15 December 2005, connection on 11 September 2020. URL : <http://journals.openedition.org/cpl/470> ; DOI : <https://doi.org/10.4000/cpl.470>

This text was automatically generated on 11 September 2020.

© All rights reserved

Implicit learning in a prediction task: Neither abstract nor based on exemplars

Arnaud Destrebecqz and Axel Cleeremans

Received March 3, 2005

Revised October 10, 2005

Accepted December 13, 2005

On line December 15, 2005

Introduction

Since about thirty-five years, research on implicit learning has shown that subjects faced with complex rule-governed situations can improve their performance without intention to learn and without clear acquisition of conscious knowledge of the rules (Cleeremans, Destrebecqz, & Boyer, 1998). Current research on implicit learning relies mainly on the artificial grammar learning (AGL) and on the sequence learning (SL) paradigms. AGL experiments have shown that participants, after being confronted with strings of letters generated by an artificial grammar, are able to identify correctly new strings as grammatical or not, in spite of the fact that they were not able to describe the rules of the grammar (Reber, 1989). In SL studies, participants perform a serial reaction time (SRT) task in which they have to indicate as fast and as accurately as possible the location of a target on a computer screen. Unknown to the participants, the sequence of stimuli follows some regularity. Typically, participants become sensitive to these regularities even though they remain often unable to access this knowledge explicitly (Destrebecqz & Cleeremans, 2001; Jiménez, Méndez, & Cleeremans, 1996; Nissen & Bullemer, 1987).

Despite numerous studies, these results are the object of ongoing controversies about what is learned when people do not know that they are learning. More specifically, several experiments have been conducted to assess whether participants acquire abstract knowledge of the rules or if it is based on memory of the training material? Three main positions about this issue have been expressed in the literature. According

to a first conception (Lewicki, Czyzewska, & Hoffman, 1987; Reber, 1989), implicit learning results in abstract knowledge, representative of the structure of the material and independent of the physical features of the stimuli. For other authors, learning is based on the memorization of fragments of the stimuli presented to the subjects (Meulemans & van der Linden, 1997; Perruchet & Amorim, 1992; Perruchet & Pacteau, 1990; Servan-Schreiber & Anderson, 1990). A third position assumes that, in the context of an artificial grammar study, participants classify new strings based on their similarity with entire training exemplars stored in memory (Brooks & Vokey, 1991). The two latter assumptions contrast with the abstractionist standpoint. Indeed, according to these hypotheses, the representations developed during learning are distributed over several memory traces and tied to the surface features of the stimuli.

In artificial grammar learning studies, a transfer procedure has been frequently used to investigate the extent to which knowledge acquired implicitly is rule-like or based on memory. In a typical transfer task, participants are presented with new strings produced with the same set of generative rules but made up of a different set of letters or symbols than the study strings (e.g., Altmann, Dienes, & Goode, 1995; Reber, 1969). The rationale underlying this procedure is that if learning reflects the abstract structure of the grammar, participants should be able to correctly classify new strings without a major cost in performance. By contrast, an important drop in accuracy is expected if learning is essentially based on memory of the training exemplars.

Significant transfer effects have been repeatedly found when different set of letters were used in training and transfer phases or when the transfer material consisted of tones, color patches, syllables or abstract symbols (Altmann et al., 1995). Even though these successful transfer effects would suggest that implicit learning is based on abstract knowledge, close inspection of classification performance revealed that a single cue, such as the identity of the initial element or illegal repetitions of the same element, was systematically used by participants to reject the ungrammatical transfer items (Tunney & Altmann, 1999). This result suggests that the transfer effect observed in previous artificial grammar learning studies cannot be attributed to the implicit abstraction of the sequential dependencies between the different elements. As previously pointed out by Tunney & Altmann, implicit grammar learning studies have shown, however, that at least some knowledge can be transferred across training and transfer material; the question at hand is to determine exactly what features in order to clarify the nature of the mechanisms subtending learning and transfer performance.

An interesting way to better understand the nature of the knowledge acquired during implicit learning episodes consists in exploring the kind of representations developed in different experimental settings. In this study, we addressed this issue using a sequential prediction task initially described by Kushner, Cleeremans, & Reber (1991).

A sequential prediction task

In the Kushner, Cleeremans and Reber (1991) experiment, participants were exposed to sequences of five stimuli presented successively on a computer screen. The task was to predict the location of the sixth stimulus. There were three possible locations (0, 1, and 2) arranged as the vertices of an invisible inverted triangle. The first five stimuli appeared at random locations but the location of the sixth stimulus depended on the spatial relationship between the second and fourth stimuli. If they appeared in the same position, then the sixth stimulus always appeared in location A (one of the three screen locations; the correspondence between sequential transitions and screen

locations was balanced between participants). If they were in a clockwise relationship, the sixth stimulus appeared in screen location B, and if they were in a counter-clockwise relationship, the sixth stimulus always appeared in the third screen location C. Only the second and fourth elements were relevant for the prediction task; the first, third and fifth stimuli were always irrelevant. This task is particularly complex because (1) there are more irrelevant than relevant events, and (2) because the location of the sixth stimuli depends on the *relationship* between the relevant stimuli, the location of each element is in itself uninformative. Despite this extreme complexity, the results of the Kushner et al. study showed that subjects became increasingly better at making accurate predictions over the 2430 trials of training, and reached at the end of the training phase a level of performance about 45% of correct responses, significantly above chance level (33%). In a second phase of this experiment, the rules were modified so that every sequence that ended in one location in the training phase now ended in another location in the transfer phase. Accuracy dropped to chance level in the transfer phase, but there was again a significant improvement in performance over the next sessions. This result showed that participants were able to transfer relatively easily from one set of sequential dependencies to another one. By contrast, in a third and final phase of the experiment, in which the sixth location was chosen at random, performance remained low and did not differ from chance level. Subjects could in some specific cases (e.g. salient sequences such as 00000 or 01010) rely on explicit knowledge to determine their response. But the authors showed that this fragmentary knowledge was clearly insufficient to account for the global level of performance. Moreover, participants were not able to describe the rules or even to differentiate the pertinent from the non-pertinent elements. The results of this experiment seem therefore to be in favor of the abstractionist standpoint according to which participants acquired implicitly rule-like knowledge about the sequential contingencies present in the training material.

In a replication of this study, Perruchet (1994) argued, however, that the increase in correct predictions was due to the memorization of the specific training sequences and not to the abstraction of the sequential rules. In this experiment, in order to pit the two hypothesis against each other, only the sequences that comprised two of the three possible instantiations of each rule (e.g. 0-0 and 1-1 as second and fourth event) were displayed during training, and the sequences comprising the remaining possibilities (e.g. 2-2) were shown in a subsequent transfer phase without feedback. The rationale underlying this procedure was that participants should respond A to transfer sequences including the pertinent combination 2-2 if they abstracted the sequential rules but they were expected to predict B or C if they simply memorized the training sequences. The transfer sequences including the combination 2-2 had indeed one additional element in common with the training sequences ending by B or C than with the training sequences ending by A. As the results of his experiment confirmed this second, memory-based, hypothesis, Perruchet argued that subjects did not abstract any rule at all, consciously or unconsciously, but that they respond to the new items based on similarity with the training exemplars.

While these results clearly show that memorization plays an important role in this prediction task, they do not rule out the possibility that some abstraction processes also subtend performance in this situation. Based on simulation results, that we briefly describe in the next section, Cleeremans (1994) suggested that participants may learn implicitly to differentiate the pertinent elements from the non-pertinent ones and

nevertheless be influenced by similarity in the transfer phase. This form of learning would constitute an abstraction in the sense that this knowledge reflects a relevant structural property of the training material.

Mechanisms

In their paper, Kushner, et al. (see also Cleeremans, 1994) showed that a buffer network (see Figure 4) was able to simulate their data. In this model, a spatial metaphor is used to represent the successive events of the sequences. Namely, five identical pools of three input units were used to represent the five elements of the sequences, each occurring in one of the three possible locations. Each input unit was connected with every hidden unit. Three output units received input from the hidden level and represented the three possible successors of the first five elements. As a simplification, elements 1 to 5 were presented at the same time at the input level. According to this procedure, the prediction task would be more akin to a categorization task in which participants have to classify the sequences in three categories based on some structural features. On each trial, the error was measured at the output level and the connection weights were modified through the back-propagation learning algorithm (Rumelhart, Hinton, & Williams, 1986).

Using this procedure, the buffer model was able to simulate participants' performance in the three phases of the Kushner et al. experiment. Learning was slower in the model at the beginning of training but this discrepancy could be attributed to the rapid memorization of salient sequences in participants — a phenomenon that the learning mechanisms instantiated in the buffer model cannot account for. Cleeremans (1994) showed that the same model was also able to simulate participants' performance in the Perruchet (1994) experiment. As for the participants, the buffer network did not abstract the generation rules of the material and there were clear indications that classification of the transfer sequence was based on similarity to stored exemplars. Cleeremans (1994) reported, however, that the representations developed by the network went beyond rote memory of the training sequences. The pattern of connection weights between input and hidden units indeed revealed that the network progressively learned to ignore information presented on the pools corresponding to non-pertinent elements. All the corresponding connection weights were very close to zero by the end of training. By contrast, the connections between the pertinent pools of input units and the hidden layer grew larger and larger during training. This result suggests that the buffer network can learn to differentiate between relevant and irrelevant sequence elements. Cleeremans (1994) also mentioned that the representations developed by the buffer network makes it possible for the model to exhibit perfect transfer with new sequences that differ from the training material only with respect to the non-pertinent elements (and not to the pertinent elements as it was the case in the Perruchet' study).

In this study, we test this prediction experimentally. To do so, we compared participants' performance in two transfer conditions differing by the nature of the transfer material. In one condition, training and transfer sequences differed by the relevant elements (as in the Perruchet' study), and by the irrelevant elements in the other condition. If participants show preserved transfer in this latter condition (as the buffer model does), it would be a good indication that they learned to differentiate between relevant and irrelevant items.

Experiment

Pilot experiments have shown that sequence learning was difficult to replicate with the original material designed by Kushner et al. We therefore simplify this material in order to promote learning. The simplification consisted in suppressing the third element of each sequence. The crucial relation that determines the position of the fifth element is now between the second and the third stimuli whereas the first and the fourth elements are always irrelevant for the task. The location of the fifth element depends exclusively on the spatial relationship between the elements 2 and 3. There are now 81 (3^4) different sequences. The response location is based on the relationship between successive events and there are as many pertinent stimuli as non-pertinent ones.

Procedure

The experiment was run on a Macintosh computer. The display consisted of three empty circles (1 cm in diameter) arranged as the vertices of an invisible equilateral triangle. The stimuli were black circles appearing within one of the three circles.

The experiment consisted of 10 sessions of 216 trials. Participants had a small rest period in the middle of each session. The 54 training sequences were presented four times in a random order during each of the 8 training sessions. Participants were simply asked to observe the sequence of four elements (stimulus duration and inter-stimuli interval were both set to 250 ms) and were then prompted to indicate the location of the fifth stimulus by using one key of the numerical keypad (the keys '4', '5', and '6' corresponded to the three corners of the invisible triangle). Subjects had 6 seconds to enter their predictions. Two different tones were used to indicate a correct or erroneous response. In case of incorrect response, the correct location was displayed on the screen for one second before presentation of the next sequence. The percentage of correct predictions was displayed on the screen after each session.

There were two transfer sessions of 216 trials (sessions 9 and 10) during which no feedback was given on the correct location of the fifth stimulus in order to prevent learning in this phase. The 54 transfer sequences were the 27 new sequences corresponding to the combinations not displayed previously and 27, randomly selected, interspersed old sequences.

Learning was compared in 2 conditions. In the rule deletion (RD) condition, the transfer sequences included new combination of relevant items, and in the context deletion (CD) condition, the transfer sequences included new combination of irrelevant items (see Annex 1 and 2).

During a subsequent rating task, subjects were asked to rate each of the four stimuli in terms of their relevance in predicting the location of the fifth stimulus on a graded scale of 1 (not important) to 5 (very important). After the experiment, subjects were interviewed about their strategies, their hypotheses about the structure of the material and the sequences that they had possibly memorized.

Stimulus material

There were 81 (3^4) different sequences. In the RD condition, the combinations 0-0, 1-2, and 2-1 between the second and third relevant elements were not displayed during the training phase and were reserved for the transfer phase (see Annex 1). In the CD condition, the same combinations, 0-0, 1-2, and 2-1, were not presented between the irrelevant first and fourth elements during the training phase and were only displayed at transfer (see Annex 2). . This procedure ensures that the three different stimuli (0, 1

and 2) were equally frequent in the four sequence positions during the training phase. It must be noted that the material presented to RD and CD participants differed during both the training and transfer phases.

Unknown to participants, the location of the fifth stimulus could be predicted based on the spatial relationship between elements 2 and 3 (identical, clockwise or counter-clockwise). The correspondence between sequence category and correct location was balanced between participants.

Participants

Twelve participants (6 in each condition) were paid 50 € for their participation in the experiment, and could earn an additional bonus of 0.02 € for each correct prediction. Participants were told that the experiment concerned the study of predictive behavior but they were not informed about the presence of sequential regularities. All participants performed the entire experiment within 5 days.

Results Training

Figure 1 (Panel A) shows the percentage of correct responses (CR) during training and transfer phases. In both conditions, performance improves gradually up to the last training session (session 8). We can also observe that performance is improved in the rule deletion condition as compared to the context deletion condition. This difference between the two conditions appears since the first session and remains relatively stable until session 8.

These impressions were confirmed by an analysis of variance (ANOVA) performed on the proportion of CR obtained during the eight sessions of training with Practice (8 levels) as a within-subject variable and Condition (2 levels) as a between-subjects variable. This analysis revealed significant main effects of Practice [$F(7,70) = 22.587$, $MS_e = 738.752$, $p < .0001$] and Condition [$F(1,10) = 7.683$, $MS_e = 1998.101$, $p < .05$]. The Practice X Condition interaction did not reach significance ($F < .4$).

Figure 1. Real (panel A) and simulated (panel B) mean percentages of correct predictions observed in the RD and CD conditions during the eight practice sessions and the two transfer sessions. The horizontal line indicates chance level (33%). Real (panel C) and simulated (panel D) percentages of correct predictions observed for the old and new sequences presented during the transfer phase. The error bars represent standard errors of the means.

Transfer

Proportion of CR did not differ between the two transfer sessions ($F < 1.8$). Therefore, we averaged performance over sessions 9 and 10 in subsequent analysis. Inspection of Figure 1 (Panel A) suggests that the introduction of novel sequences in the transfer phase resulted in a drop in accuracy essentially in the RD condition. This impression is confirmed by an ANOVA in which we compared the mean accuracy during the two last sessions of training (session 7 and 8) with the mean accuracy during the two transfer sessions. This analysis was performed with Session (2 levels) as a within-subject variable and Condition as a between-subjects variable. The ANOVA revealed a significant main effect of Session [$F(1,10) = 18.500$, $MS_e = 617.780$, $p < .01$] and a significant Session 5 Condition interaction [$F(1,10) = 7.106$, $MS_e = 237.290$, $p < .05$]. The main effect of Condition did not reach significance ($F < .5$). Planned comparisons indicated that accuracy dropped between training and transfer phases in the RD condition [$F(1,10) = 15.435$, $MS_e = 810.410$, $p < .05$] but not in the CD condition ($F < 3.2$).

Figure 2. Mean ratings of the four sequence elements in terms of their relevance in predicting the location of the fifth trial plotted separately for the CD and RD conditions.

To further analyze transfer performance, we compared accuracy between old and new sequences presented during the two transfer sessions (see Figure 1, Panel C). We performed another ANOVA with Condition (2 levels) as a between-subjects variable and Sequence type (2 levels) as a within-subject variable. This analysis revealed a significant effect of Sequence type [$F(1,10) = 30.424$, $MS_e = 1643.415$, $p < .001$].

Image8.001] and a significant Sequence type 5 Condition interaction [$F(1,10) = 9.759$, $MS_e = 527.156$, $p < .05$]. The main effect of Condition did not reach significance ($F < .06$). Planned comparisons revealed that the main effect of

Sequence type was significant in the RD condition [$F(1,10) = 29.163$, $MS_e = 2016.058$, $p < .01$] but not in the CD condition ($F < 4$).

To summarize, these analysis indicated that performance remained stable between training and transfer phases in the CD but not in the RD condition in which we observed impaired transfer performance. In this latter condition, accuracy dropped for new sequences, as compared to old sequences, and did not differ from chance level ($t < 0.3$).

Ratings

When asked to rate each event according to their importance for the prediction task, RD and CD participants failed to differentiate between relevant and irrelevant items (see Figure 2).

Nevertheless, they often explicitly and accurately reported some particularly salient sequences. These sequences were made by three or four repetitions of the stimulus at the same location, and they all predict the same particular location for the fifth event. One RD participant and one CD participant have been able to rate relevant and irrelevant items correctly by attributing the lowest rating to elements 1 and 4 and the highest rating to elements 2 and 3. However, the RD participant was not able to describe the relationship between these relevant items and the fifth location. The CD participant could tell that when elements 2 and 3 appeared in the same location, the fifth element appeared in location A.

Simulations

How well would the buffer network (see Figure 3) learn the material used in this experiment? To find out, we trained six buffer networks with four pools of three input units, each of these pools representing the four sequence elements. Three output units were used to represent the three possible locations of the fifth element and the hidden layer was comprised of 5 units.

Each network was initialized with random weights between -0.5 and 0.5 and presented with the same material, and for the same number of trials, as participants. The task of the network was also to predict the fifth element of each sequence. We used the same simplification procedure as Cleeremans (1994) and presented the four elements simultaneously to the networks. The model was then equivalent to a three-layers backpropagation network. During the training phase, the learning rate and momentum parameters were set to 0.5 and 0.9 respectively. During the transfer phase, these parameters were set to 0.0 because participants did not receive any feedback on their accuracy during this phase of the experiment.

Figure 1 (panel B) shows the mean network performance for the training and transfer phases in the two conditions. Results indicate that the buffer network can account for the main aspects of human performance. The percentage of correct responses tends to increase with training in both conditions. As in the experiment, the introduction of novel sequences in session 9 exerts a detrimental effect on performance in the RD condition but not in the CD condition. Figure 1 (panel D) also indicates that, as in participants, networks' performance did not differ between old and new transfer sequences in the CD condition but was clearly impaired for new sequences as compared to old sequences in the RD condition.

Figure 3. The buffer network : Each pool of three input units represents one of the four sequence element. Three output units are used to represent the fifth element predicted by the network.

Image9Simulations are not perfect however and present a series of discrepancies with participants' performance. The mean level of performance is slightly underestimated by the model in the CD condition and clearly overestimated in the RD condition. In this latter condition, the simulated percentage of correct responses tends to remain relatively stable from the second to the last training session while performance's improvement is much more progressive in participants. Differences in learning curves between buffer networks and participants were previously reported by Cleeremans (1994) and may be related to the fact that participants also rely on memory for specific instances during training. Particularly salient sequences consisting for instance in the repetition of a single element or in simple alternations between two locations may be quickly learned by participants during the first training session while the buffer network has no computational mechanism to capture this aspect of performance. This might explain the initial difference between networks and participants in the CD condition.

By contrast, in the RD condition, the model quickly reaches his performance peak from session 2. We observed this pattern of results in our simulations using different parameters values: learning was systematically more important and occurred faster for networks trained in the RD condition than for those trained in the CD condition. This pattern of results reproduces, although in an emphasized way, what we observed in the experiment where performance was also improved in the RD condition. We discuss this unexpected result in the next section.

Discussion

The notion that rule-based learning can occur implicitly has been previously rejected based on the observation that abstract information was not necessary to perform the tasks used to assess the acquired knowledge (Perruchet, Gallego, & Savy, 1990; Perruchet & Pacteau, 1990). Namely, in the case of the prediction task used in this study, Perruchet (1994) has shown that transfer performance was determined by the similarity between novel and training sequences rather than by the rule-based category of the transfer sequences. Our results are in line with this assumption. Indeed, in our experiment, participants also tend to respond to transfer sequences based on their

similarity to the training sequences. For instance, in the RD condition, participants tended to erroneously respond B or C to the transfer sequence 0000 because the more similar training sequences 0010, 0200, 0020 and 0100 were followed by locations B or C. By contrast, in the CD condition, participants correctly responded A to the transfer sequence 0000 because the more similar training sequences 0001, 0002, 1000, and 2000 are also followed by location A.

In a previous study, Cleeremans (1994) claimed that successful transfer to sequences with new irrelevant contexts could indicate that participants learned to differentiate between relevant and irrelevant items. The behavioral and simulation results of this study tend to support this assumption. While we do not dispute the role of similarity in transfer performance, our results also suggest that learning in the prediction task was not exclusively based on raw memory for exemplars. Indeed, we

observed that the percentage of correct predictions was systematically higher in the RD condition than in the CD condition throughout the eight training sessions. Close inspection of the training sequences reveals that this result could be attributed to a structural difference between both sets of training sequences. In the RD condition, each stimulus 0, 1, or 2 appearing in the second or third locations (i.e., the relevant sequence items) can only be followed in the fifth location by two of the three possible stimuli. For example, if the second element appeared in location 0 then the fifth element can only appear in location B or C. By contrast, in the CD condition, any stimulus appearing in the second or third position can be followed by the three possible locations in the fifth location. As a result, there is a crucial difference between the two sets of training sequences with respect to the amount of information that can be extracted based on the relevant elements of each training sequence. In the RD condition, each relevant item conveys individually more information about the location of the fifth trial than in the CD condition. In the RD condition, the second and third elements reduce the uncertainty associated with the identity of the fifth trial. In the CD condition, the fifth location can only be predicted based on the relationship between the relevant elements. The positions of the first and fourth sequence elements are irrelevant in both conditions.

Figure 4. Luce ratio of the summed connection weights between the four pools of input units and the pool of hidden units plotted separately for the RD and CD conditions. These values denote the relative importance of each represented sequence element for the prediction task.

Simulation results indicate that the performance of the buffer network was strongly influenced by this structural difference between training sets. As indicated previously by Cleeremans (1994), the buffer model learns to ignore the information coming from pools of input units coding for irrelevant elements. As illustrated in Figure 4, at the end of the training phase, the connection weights between input and hidden units are stronger for the second and third pools of input units, coding for the relevant sequence elements, than for the first and fourth pools of input units that code for the irrelevant elements. This learning process is improved in the RD condition because the network can start to use the information provided by the second and third elements in predicting the fifth trial even before it has developed representations taking the relationship between these two elements into account. As a result, network's performance improves quickly during the first two sessions in the RD condition and remains relatively stable until the transfer phase, while it improves much more gradually in the CD condition. This can also explain why performance remains systematically lower in the CD condition as compared to the RD condition because, in this latter condition, the network keeps to benefit from the structural difference in training sets that has boosted its performance at the early stages of learning.

Participants' performance was also improved in the RD condition, as compared to the CD condition, in our experiment. This result suggests that participants were also influenced by the structural difference between the CD and RD training sets and might therefore indicate that participants learn, as the model, to differentiate between relevant and irrelevant items.

This idea seems at odds with the notion that performance improvement in this task simply reflects the increasing number of sequences memorized by the participants. Such a learning mechanism could not explain the improved performance of

participants trained in the RD condition. Indeed, memorization of the four sequence elements would allow to predict the location of the fifth trial equally well in both RD and CD conditions. Whether the ability to extract the relevant features of the sequential material constitutes abstract or rule-like knowledge is an open question, however, it certainly involves more sophisticated learning processes than rote memory of exemplars.

Did learning occur implicitly or explicitly? Only one participant in each condition was able to correctly rate the four elements according to their importance for the prediction task, and only one of them was able to accurately state one component of the prediction rule. None of the other participants was able to show any conscious knowledge of the sequential regularities. Most of them, however, were able to report some salient sequences involving repetitions of the same stimulus or alternations between two locations.

As we said earlier, the buffer network has no computational mechanism to account for this sensibility to salient sequences, which undoubtedly influence prediction performance. An important contribution of this study, however, is to demonstrate that the associative learning mechanisms implemented by the buffer model are able to account for both the sensitivity to the relevant features of the material and for the influence of similarity to training exemplars.

Different putative mechanisms have been proposed to account for performance in implicit learning studies: rule abstraction, memorization of training instances, and sensitivity to the statistical properties of the environment. The results of this study suggest that performance in this prediction task is based on learning processes, such as those implemented in the buffer network, resulting in the acquisition of graded and distributed knowledge. Learning, in this perspective, consists in the development of an increased sensitivity to the most relevant source of information and depends on the structural properties of the training environment (see also Gomez, 2002). In our view, the abstract nature of the knowledge acquired during a learning episode evolves along a graded dimension going from simple memorization to rule abstraction. Connectionist modeling makes it possible to go beyond descriptive theories and to identify the nature of the representations developed throughout learning.

Acknowledgements

AD and AC are respectively Scientific Research Worker and Senior Research Associate of the National Fund for Scientific Research (FNRS, Belgium). AD is now supported by a post-doctoral grant from the Fyssen Foundation.

Thanks to Pawel Lewicki and an anonymous reviewer for their useful comments to a previous version of this article.

Annex 1.

This table shows the sequences presented in the RD condition. The framed sequences were presented in the transfer phase.

Same location Response A	Clockwise relationship Response B	Counter-clockwise relationship Response C
0000	0010	0020
0001	0011	0021
0002	0012	0022
1000	1010	1020
1001	1011	1021
1002	1012	1022
2000	2010	2020
2001	2011	2021
2002	2012	2022
0110	0120	0100
0111	0121	0101
0112	0122	0102
1110	1120	1100
1111	1121	1101
1112	1122	1102
2110	2120	2100
2111	2121	2101
2112	2122	2102
0220	0200	0210
0221	0201	0211
0222	0202	0212
1220	1200	1210
1221	1201	1211
1222	1202	1212
2220	2200	2210
2221	2201	2211
2222	2202	2212

Annex 2

This table shows the sequences presented in the CD condition. The framed sequences were presented in the transfer phase.

Same location Response A	Clockwise relationship Response B	Counter-clockwise relationship Response C
0000	0010	0020
0001	0011	0021
0002	0012	0022
1000	1010	1020
1001	1011	1021
1002	1012	1022
2000	2010	2020
2001	2011	2021
2002	2012	2022
0110	0120	0100
0111	0121	0101
0112	0122	0102
1110	1120	1100
1111	1121	1101
1112	1122	1102
2110	2120	2100
2111	2121	2101
2112	2122	2102
0220	0200	0210
0221	0201	0211
0222	0202	0212
1220	1200	1210
1221	1201	1211
1222	1202	1212
2220	2200	2210
2221	2201	2211
2222	2202	2212

BIBLIOGRAPHY

- Altmann, G. T. M., Dienes, Z., & Goode, A. (1995). On the modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology-Learning Memory and Cognition*, 21, 899-912.
- Brooks, L. R., & Vokey, J. R. (1991). Abstract Analogies and Abstracted Grammars - Comments on Reber (1989) and Mathews Et-Al (1989). *Journal of Experimental Psychology-General*, 120(3), 316-323.
- Cleeremans, A. (1994). The representation of structure in sequence prediction tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV : Conscious and unconscious information processing* (pp. 783-810). Cambridge: MIT Press.
- Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: news from the front. *Trends in Cognitive Sciences*, 2(10), 406-416.
- Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin & Review*, 8(2), 343-350.
- Gomez, R. (2002). Variability and detection of invariant structure. *Psychological Science*, 13, 431-436.
- Jiménez, L., Méndez, C., & Cleeremans, A. (1996). Comparing direct and indirect measure of sequence learning. *Journal of Experimental Psychology-Learning Memory and Cognition*, 22, 948-969.
- Kushner, M., Cleeremans, A., & Reber, A. S. (1991). *Implicit detection of event interdependencies and a PDP model of the process*. Paper presented at the Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society.
- Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology-Learning Memory and Cognition*, 13, 523-530.
- Meulemans, T., & van der Linden, M. (1997). The associative chunk strength in artificial grammar learning. *Journal of Experimental Psychology-Learning Memory and Cognition*, 23, 1007-1028.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirement of learning : Evidence from performance measures. *Cognitive Psychology*, 19, 1-32.
- Perruchet, P. (1994). *Learning from complex rule-governed environments: On the proper functions of nonconscious and conscious processes*. Paper presented at the Attention and Performance XV: Conscious and nonconscious information.
- Perruchet, P., & Amorim, M. A. (1992). Conscious knowledge and changes in performance in sequence learning : Evidence against dissociation. *Journal of Experimental Psychology-Learning Memory and Cognition*, 18, 785-800.
- Perruchet, P., Gallego, J., & Savy, I. (1990). A critical reappraisal of the evidence for unconscious abstraction of deterministic rules in complex experimental situations. *Cognitive Psychology*, 22, 493-516.
- Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning : Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology : General*, 119, 264-275.
- Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. *Journal of Experimental Psychology*, 81, 115-119.

- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology : General*, 118, 219-235.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (pp. 318-362). Cambridge: MIT Press.
- Servan-Schreiber, E., & Anderson, J. R. (1990). Learning Artificial Grammars with Competitive Chunking. *Journal of Experimental Psychology-Learning Memory and Cognition*, 16(4), 592-608.
- Tunney, R. J., & Altmann, G. T. M. (1999). The transfer effect in artificial grammar learning: Reappraising the evidence on the transfer of sequential dependencies. *Journal of Experimental Psychology-Learning Memory and Cognition*, 25(5), 1322-1333.

ABSTRACTS

The notion that rule-based learning can occur implicitly has been previously challenged based on the observation that abstract information was not always necessary to perform the tasks used to assess the acquired knowledge. Some authors suggest instead that implicit learning is based on memorization of training material. In this study, we address this issue in the context of a sequential prediction task initially described in Kushner, Cleeremans & Reber (1991). The task consists in predicting the location of the fifth element of a sequence amongst three possible locations. Unknown to the participants, the correct location can be predicted based on the relationship between two of the four preceding sequence elements. After training, we compared transfer performance in two conditions. In the “rule deletion” condition, transfer sequences contained new combinations of relevant elements and in the “context deletion” condition, new combinations of irrelevant elements. Based on behavioral and modeling results, we confirm the strong influence of similarity in transfer performance but, crucially, we also conclude that participants progressively learned implicitly to differentiate between relevant and irrelevant elements for the prediction task — a learning process that is not equivalent to rule abstraction but that is clearly a step away from rote memorization.

L'hypothèse selon laquelle l'apprentissage de règles peut avoir lieu implicitement a été remise en question car, dans plusieurs études, des connaissances abstraites se sont révélées inutiles pour accomplir les tâches utilisées pour mesurer l'apprentissage. Dans cette étude, nous étudions cette question à l'aide d'une tâche séquentielle de prédiction décrite initialement par Kushner, Reber, et Cleeremans (1991). La tâche consiste à prédire la position du cinquième élément d'une séquence parmi trois positions possibles. La réponse dépend, à l'insu des sujets, de la relation existant entre deux des quatre éléments de la séquence. Après l'apprentissage, la performance de transfert est comparée dans deux conditions. Dans la condition “Rule Deletion”, les séquences de transfert incluent de nouvelles combinaisons d'éléments pertinents et, dans la condition “Context Deletion”, de nouvelles combinaisons d'éléments non-pertinents. Sur base des résultats comportementaux et de simulations connexionnistes, nous confirmons l'influence importante de la similarité dans le transfert mais nous montrons également que les sujets ont appris implicitement à différencier les éléments pertinents et non-pertinents — un processus d'apprentissage qui ne peut être assimilé à de la simple mémorisation.

INDEX

Keywords: implicit learning, Abstraction, Connectionist Modeling

AUTHORS

ARNAUD DESTREBECQZ

Cognitive Science Research Unit Université Libre de Bruxelles 50 av. FD Roosevelt (CP 191)1050 Brussels BELGIUM

AXEL CLEEREMANS

Cognitive Science Research Unit Université Libre de Bruxelles