

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Hebbian Learning of Artificial Grammars

Permalink

<https://escholarship.org/uc/item/9q32w0xj>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 14(0)

Authors

Ganis, Giorgio

Schendan, Haline

Publication Date

1992

Peer reviewed

Hebbian Learning of Artificial Grammars

Giorgio Ganis

Department of Cognitive Science
9500 Gilman Dr., 0515
UC, San Diego
La Jolla, CA 92093-0515
ganis@cogsci.ucsd.edu

Haline Schendan

Program in Neurosciences, 0608
UC, San Diego
La Jolla, CA 92093
hschenda@igrad1.ucsd.edu

Abstract*

A connectionist model is presented that used a hebbian learning rule to acquire knowledge about an artificial grammar (AG). The validity of the model was evaluated by the simulation of two classic experiments from the AG learning literature. The first experiment showed that human subjects were significantly better at learning to recall a set of strings generated by an AG, than by a random, process. The model shows the same pattern of performance. The second experiment showed that human subjects were able to generalize the knowledge they acquired during AG learning to novel strings generated by the same grammar. The model is also capable of generalization, and the percentage of errors made by human subjects and by the model are qualitatively and quantitatively very similar.

Overall, the model suggests that hebbian learning is a viable candidate for the mechanism by which human subjects become sensitive to the regularities present in AG's. From the perspective of computational neuroscience, the implications of the model for implicit learning theory, as well as what the model may suggest about the relationship between implicit and explicit memory, are discussed.

Introduction

An artificial grammar (AG) is produced according to a "Markovian process in which a transition from a state S_i to any state S_j produces a letter" (Reber, 1967). Experimental psychology has extensively studied the learning of AG's (Reber, 1989). The data reveal several things. I) Subjects can learn and use the grammatical structure of the strings to facilitate the learning itself. This has been shown by comparing the learning time for grammatical and random letter strings. II) Subjects can generalize the knowledge they have acquired to novel strings. This was found by using a discrimination task in which previously unseen strings were classified as

grammatical or non-grammatical. III) Depending upon the instructions, subjects have differing degrees of awareness of the nature of the grammar and of the underlying rules. Nevertheless, both the processes involved in AG learning and the resultant knowledge structure remain largely unknown. The purpose of our model was to shed light upon both of these aspects of AG learning.

There are several ways in which AG knowledge may be represented. I) Subjects could represent the AG by explicit rules that correspond to the Markovian process. II) Subjects may discover the correlation between symbols, combinations of symbols, and their string positions, as well as the frequency of occurrence of such correlations. III) According to the "distributive position" (Vokey and Brooks, 1991), the knowledge that subjects acquire is composed of a set of specific instances. These instances can be used to categorize new stimuli by means of a similarity evaluation process. IV) In contrast to III, the "abstractive position" concludes that subjects learn by means of a nonconscious abstraction system without retention of specific instances (Reber, 1989). According to this position, both learning and knowledge are implicit.

While there is evidence for all of these positions, this may not be problematic because there may be multiple learning mechanisms which thus result in multiple forms of knowledge representation. Additionally, these positions may not necessarily be mutually exclusive. For example, the distributive position bears some resemblance to exemplar models of perceptual categorization and the abstractive position resembles prototype models (Posner and Keele, 1968). Parallel distributed processing (PDP) models have been proposed which result in representations that encompass both exemplars and prototypes (Rumelhart and McClelland, 1986). PDP models are an attempt to use what is known about the brain to constrain psychological explanation. With regard to the kind of sequence learning that may be involved in AG learning, Cleeremans and McClelland (1991) used a simple recurrent PDP model to explain how human subjects could encode the temporal context of complex sequential material. They suggest that "sequence learning may be based solely on associative learning processes." The model presented here suggests that

* Authors names are in alphabetical order. Research supported by the McDonnell-Pew Center for Cognitive Neuroscience in San Diego.

associative learning, as may be involved in AG learning, may take place according to a hebbian mechanism.

Though not a PDP model, the competitive chunking model of Servan- Schreiber and Anderson (1990), one of the few computational models of AG learning, is also of interest because it is based on the notion of hierarchical chunking. Chunks within a level compete with each other, while those between levels do not. The main shortcoming of the model, however, is that it suggests little about neuronal mechanisms, failing to explain how the mind/brain system can accomplish AG learning by chunking.

One way in which chunks may be formed is through a process that causes multiple representations to be connected together. The mechanism proposed by our model is that of hebbian modification of synaptic weights between neurons in a network. The Hebb rule is a neurobiologically plausible mechanism of learning (for review, cf. Brown et al, 1990). It states that when cell A excites cell B "some growth or metabolic" change will take place in both cells such that there is an increase in the ability of cell A to fire cell B: (Hebb, 1949). The main points of hebbian plasticity are that (1) it is activity between two neurons which determines plasticity, (2) the level of activity in cell B must be high enough to generate action potentials, and (3) the change in efficacy is specific to the connection between cell A and B. In other words, correlated firing in two neurons enables an increase in efficacy of the connection between them.

One of the hallmarks of hebbian plasticity is that it is a mechanism whereby correlations in the environment are the parameters which determine learning and memory. Correlations between environmental stimuli are a ubiquitous characteristic of both the structure of the world and the temporal sequencing of events. The PDP model presented in this paper incorporates the Hebb rule since it is the appropriate computational function for picking up the correlations present in AG strings. In this paper, we focus upon the computations and products of neuronal plasticity which enable a rapprochement between the various theories of AG learning. The model is tested on the seminal learning experiments conducted by Reber (1967). Additionally, our model may have general implications for how the brain acquires and structures its knowledge of the world.

The model

The model is a feedforward net composed of three layers, as shown in Figure 1. The input layer is a 6X8 two dimensional array in which the rows correspond to the set of symbols that may appear in a string plus a *blank* symbol, and the columns correspond to the position of the symbol within each string. Each unit projects forward onto several units in the next layer. The dimensions of layers 2 and 3 are 10X8. Each unit in layers 2 and 3 receives projections from all units in the previous layer corresponding to its position and to positions immediately adjacent to it.




Figure 1. Architecture of the network. Only a few connections are depicted for simplicity. The filled ellipses represent the units in the input layer that are activated by the string -bottom.

The training process can be divided into stages. First, a string was presented and the activation of each unit was computed in a feedforward fashion by taking a weighted sum of the incoming inputs. The formula used was:

$$(1) \text{Output}_i = \beta \sum_j w_{ij} \text{Input}_j, \text{ where } \beta \text{ is a scaling factor.}$$

In order to prevent the activations from becoming too high, they were clipped at top and bottom. Then, for each position, the most active unit was selected and its weights updated according to a hebbian rule:

$$(2) w = \text{Input}_j \text{Output}_i \text{Delta, where Delta is the learning rate.}$$

After being updated, the weights were normalized in such a way that the sum of their absolute value remained constant. Such a normalization can be justified in neurobiological terms by assuming that weights have a metabolic cost; there is a limit to the amount of connection weighting that a neuron may support. This process is repeated for every string and the changes in the weights are cumulative. The value of Delta was selected following pilot simulations and analysis of the developed weights and activations. Delta of 0.2 was chosen because this value resulted in more local weight representations, while allowing more units to develop than did higher values. Lower values of Delta resulted in more distributed weight representations.

Experiment 1

The first experiment simulated experiment 1 of Reber (1967). In this experiment there were two types of 6-8 letter strings, generated either by a Markovian (grammatical) or a random process. Fig. 2 shows the




Figure 2. The artificial grammar used both in Reber (1967) and in our experiments.

Markovian process that was used both in our experiments and in Reber (1967). Subjects were trained on recall of 4 letter strings. For each of 7 sets of strings, training continued until a criterion per set was reached. Grammatical strings were learned faster than random, as reflected in the mean number of recall errors per set (ME). Supposedly, the acquisition of knowledge about the AG made it easier to recall, and/ or learn to recall, the grammatical strings. Analogously, the model was trained on either grammatical or random strings, and the training list was divided into sets of 4 strings. The network learned by updating its weights following each presentation of a string. To evaluate performance of the net, it was necessary to select a criterion that was analogous to ME. Average mean activation (avgMA) per set, which is inversely related to ME, was chosen because it incorporates the values of the weights that develop during training. AvgMA was calculated for both grammatical and random strings for 4 subjects.

Results

The graphs of the model's data for both layers 1 (Fig. 3a) and 2 (Fig. 3b) approximate Reber (1967) well. For both layers 1 and 2, the graph for grammatical strings rises with each set, while that for the random strings remains relatively flat. Performance improved with training set and grammatical strings were easier to learn than random. A two-way ANOVA was performed. For layer 1, there was a significant increase in avgMA across sets for both grammatical and random strings [(F = 72.78), $p < 0.001$]. String type, grammatical versus random, was also significant [(F = 225.37), $p < 0.001$]. Additionally, there was an interaction effect of string type with set number [(F = 29.33), $p < 0.001$]. For layer 2 also, there was a significant increase across training sets for both grammatical and random strings [(F = 70.18), $p < 0.001$]. String type, grammatical versus random, was also significant [(F = 306.95), $p < 0.001$]. Additionally, there was an interaction effect of string type with set number [(F = 22.63), $p < 0.001$]. A Tukey test for the avgMA for




Figure 3a. Results of Experiment 1 for layer 1.




Figure 3b. Results of Experiment 1 for layer 2.

sets 1 and 2 was not significant for both layers 1 and 2. This accounts for the interaction. The most striking difference between the experimental and simulation graphs corresponds to the steep change that Reber found between sets 1 and 2 for both string types. This is missing from both simulation graphs.

Discussion

The absence of a steep rising phase during training on sets 1 and 2 is consistent with Reber's explanation of this finding. He attributed this rise to "a rather complicated warm-up effect" (Reber, 1967). This learning is unrelated to the stimuli. Rather, it is a kind of procedural learning of the training task, irrespective of the stimuli. It is of interest that procedural learning is also a kind of implicit learning (Squire, 1987). Thus there may be at least two forms of implicit learning active that are involved in AG learning. One that is stimulus specific and one that is task general. It is only the former which interested Reber and which the model was designed to emulate.

Experiment 2

Experiment 2 of Reber (1967) was also simulated. Here, subjects were trained with 20 grammatical letter strings until a criterion level of performance was reached. At test, subjects classified novel strings as grammatical or ungrammatical. These grammaticality judgments were about 79% correct (Reber, 1967).

While the simulation training procedure was identical to that used in experiment 1, the training strings differed as per Reber (1967). 20, as opposed to 28, grammatical strings were learned, and the range of string lengths was broader, 3-8 symbols long. Each sequence of 20 grammatical strings was presented to the net 5 times, for a total of 100 training trials. The testing phase and the construction of grammatical, ungrammatical, and random test strings were identical to that of Reber (1967). The trained net was presented with all 44 test strings twice. 5 subjects were simulated by using different initial weights. In experiment 1, avgMA was used as the dependent variable because it was a rough index of learning; in experiment 2, fine grammaticality judgments were required, based upon subtle irregularities present in the strings. We therefore used a more sensitive criterion for grammaticality judgment. The criterion was such that, after being trained with grammatical strings, the net classified such strings as grammatical. At the end of training, the net was tested with its training strings. Each string generated different activations in the various units. For each one of the eight possible spatial positions, the minimum activation obtained over the whole set of strings was taken as criterion for that spatial position. A string was judged as grammatical by a layer if and only if the activations generated by that string were above criterion for all the eight possible spatial positions. If there was a conflict between the judgments made by the two layers, the choice was made at random. It is obvious by construction that all the training strings were judged grammatical.

Results

The frequencies of errors made by the network matched those of Reber (1967), as shown in Table 1a and 1b, respectively. The type of errors made by the network on ungrammatical items was analyzed. As in Reber (1967), strings with multiple errors had a significantly higher detection rate than any of the others, [$X^2(1) = 6.21$, $p < 0.01$], the detection rate of strings with a single error in the last position was higher than that of strings with an error in the middle, [$X^2(1) = 6.42$, $p < 0.01$], and the detection rate of strings with a single error in the first position was higher than that of strings with an error in the last position, [$X^2(1) = 6.94$, $p < 0.01$]. The analysis of the weight matrices indicated that n-gram detectors had developed. N varied from 1-3 symbols for units in layer 1; this means that the units could be selective for a single symbol, bigram, or trigram. An average of 40% of the units in layer 1 became selective for a single n-gram. 5% of the units became selective for

two n-grams. The remaining units had very small weights, non-selective for any particular n-gram. An example of a single symbol detector is a unit which was highly activated only by a "4" in the second position. An example of a single bigram detector is a unit which became selective only for the bigram "44" in positions 3 and 4. Examples of trigram detectors are a unit selective for the trigram "200" in positions 1, 2, and 3, and a unit selective for the trigram "430" in positions 4, 5, and 6. Units in layer 2 were always highly activated by multiple symbols, bigrams, trigrams, and/or higher n-grams.

Discussion

The consistency of the model with Reber's data support the hypothesis that the learning principle of the model is analogous to that used by human beings. Assuming this, then the structure of the weight matrices developed according to this learning principle may reveal something of the nature of mental/brain representation of AG knowledge. While Reber (1967) found that errors in the last position were most consistently detected, our model was most consistent at detecting errors in the first position. This discrepancy between Reber (1967) and the model arose because there was less variability that position than in any other position. Specifically, while the last symbol in a string was always a '1' and the first symbol either a '2' or '3', the position of the final symbol could appear in positions 3-8, whereas the initial symbol could only appear in position 1. Thus, considering both variability in symbol type and position, symbols in the first position were most determined. Additionally, units in the first and last positions received input from only 2 columns of units in the previous layer, whereas all other units received input from 3 columns, and the constraint on maximum total weight yielded higher weights per connection between units for those receiving 2 columns of input. Since symbols at the end of a string could be represented by a unit in positions 3-8, strings with the final symbol in either of positions 3-7 would have lower connection strengths than for symbols in either the first or 8th position. Thus the end symbols were disadvantaged, relative to the those in the first position, in terms of the strength of weights that may project to them. This further contributed to the superior error detection for symbols in the first position.

		Gr	Ungr
		79%	21%
Judged	Gr	79%	21%
	Ungr	22%	78%

Table 1a. Frequencies of errors obtained in Reber (1967).

		Gr	Ungr
Judged	Gr	73%	27%
	Ungr	23%	77%

Table 1b. Frequencies of errors obtained in the model.

The analysis of the weight matrices confirmed that pooling knowledge present in both layers 1 and 2 is a good strategy. Both layers developed units that preferred certain n-grams to others. However, there was a broad range of n-gram specificity in each layer. Some units were specific to only one n-gram, while others responded well to many. In general, this specificity reflected two characteristics of the AG used. One characteristic is the frequency of occurrence of each n-gram. More frequent n-grams had a higher probability to develop responsive units, and those units were also more likely to respond only to that n-gram. The second characteristic is number of different symbols or n-grams which may occur per position, or set of positions, respectively. The more n-grams, the less specificity there was for any particular n-gram at that position.

One problem with the representation of n-grams within layers 1 and 2 needs to be emphasized. Because of the two characteristics of the AG mentioned above, the broad tuning of some units tended to make such units more highly activated by some ungrammatical n-grams. These ungrammatical n-grams were mixtures of the multiple, grammatical n-grams that the unit had come to prefer. For example, a unit that preferred legal n-grams "444" and "002" in positions 2, 3, and 4 could combine these in several ways that result in a strong activation for ungrammatical n-grams, such as "404." This property of broadly tuned units indicates that for AG learning, distributed representation can increase error rates. Indeed, several of the errors made in grammaticality judgments were the results of this. Had there been more unit specificity, the accuracy of the model might have been greater. However, since the model performed about as well as Reber's subjects, humans may also use broadly tuned units to code AG information. Thus the successes and pitfalls of hebbian learning mirror those of implicit learning mechanisms in humans.

It should also be mentioned that layer 2 units were more broadly tuned than layer 1 units. This follows naturally from the fact that layer 2 made use of the units in layer 1, many of which themselves were broadly tuned. That the combined information from both layers yielded greater accuracy is also of relevance to theories about brain processing. According to neuroscience, it is the more advanced the levels of processing are the ones that most strongly influence overt behavior. Assuming the model captures the most relevant aspects of the neurobiology of learning, then the simulation results

argue against this idea. It is perhaps more realistic to say that the level of processing that is most relevant to the overt response depends upon the characteristics of the task. In the case of AG learning, each layer contributes different factors to the grammaticality judgments. Layer 2 has units which prefer longer strings than any layer 1 unit. For strings, or n-grams, less than six symbols long, there may be a unit in layer 1 that prefers it to any other string. Thus, the most important function of layer 2 is to represent longer strings (than layer 1). The most important function of layer 1, however, may be to be specific to only one n-gram. This may counteract the effects of the broadness of layer 2 units. Thus it is a combination of narrow and broad tuning in conjunction with distributed representation that constitutes the useful structure of AG knowledge.

General Discussion

In this paper, a PDP model was presented that used a hebbian learning rule to acquire AG knowledge. To evaluate the validity of the model, two experiments from the classic paper by Reber (1967) were selected. It was found that the model simulated the Reber results well. The increase in avgMA of the net paralleled the decrease in ME that humans exhibit. Novel stimuli that display the same invariants as learned stimuli elicit similar responses both in the model and in human subjects. Thus both are capable of generalization. This was demonstrated by the testing phase of experiment 2. Thus, in general, a hebbian learning mechanism is capable of computing the correlations between aspects of a stimulus and is able to use that knowledge to generalize and make decisions. These are basic cognitive abilities.

The model may have important implications extending beyond the domain of AG learning. While the model may be presumed to reveal characteristics only of implicit learning, our results are consistent with the hypothesis that implicit learning is carried out in the brain by means of a hebbian mechanism. Perhaps, the representations formed consist of units distributed across areas of cortex, with both narrow and broad tuning for correlations in the stimuli.

This characterization of the structure of knowledge resulting from implicit learning may be used to address issues from the implicit/explicit memory literature, such as whether or not there is a relationship between implicit and explicit stores and, if so, what the nature of the interaction may be. Specifically, a grammaticality judgment must be verbalized. Thus implicit knowledge must be accessible to verbalization, or explicit memory, systems. The simulations of Reber's experiments suggest how implicit knowledge may be used by the explicit system.

Once the model has been trained on a set of AG strings, within it is stored the level of activation that corresponded to acceptable sequences. This information exists as the weights connecting units in the model and, perhaps, as the relative efficacy of synapses in the brain of a human

subject. In this type of neural representation, units coding grammatical n-grams have higher weights for symbols in correct positions within a string. Thus when a novel string is presented, mainly grammatical n-gram units will be activated. These grammatical n-gram units will produce higher activation over the net than when ungrammatical strings are presented. In the ungrammatical case, fewer of these highly weighted n-gram units are activated, resulting in relatively lower net activation for such strings.

Through training, systems involved in generating a response may learn what level of activation tends to be elicited by the training strings. Any novel string that elicits a similar level of net activation, its minimum determined by a threshold, may tend to be treated in the same way by the response system. This may be related to perceptual categorization (Posner and Keele, 1968).

That this may be a consequence of AG learning is supported by the Reber et al (1980) finding that the earlier explicit instructions are presented to a subject during training, the better will be performance when tested with a well formedness task. Well formedness was substituted for grammaticality by Reber in subsequent experiments (Reber, 1989). The implications of this study for the results of the model are twofold. One, the function of the explicit instructions may be to provide the learning system with feedback. This may enhance the ability of the system to categorize the stimuli as belonging to one kind. Additionally, this may enable a response system to better differentiate between the levels of activations that correspond to the category of training stimuli. For example, in terms of the simulation, the threshold, or range of thresholds, may be set more precisely.

The second implication of the Reber et al (1980) results for the model is that the explicit system is capable of interaction with implicit knowledge systems. The nature of this interaction probably depends upon the specific task demands. However, it is not clear in what ways they interact. The explicit system may not interact with the kind of learning system modelled here, but rather with a system involved in generating responses. However, since in the Reber experiments, training did not involve grammaticality feedback, the human learning system can acquire the ability to make accurate well formedness judgments without being specifically taught to do so. Rather, human subjects can use the knowledge gleaned from exemplar presentation to influence such judgments. Our model supports this view. The hebbian learning mechanism permits knowledge of correlations within exemplars to be stored in a form that may be used by an explicit, verbalization, or response system to successfully complete tasks like grammaticality or well formedness judgments. The model suggests also that the learning system may function largely in isolation from one involved in determining the overt behavioral response. However, that human subjects are able to use explicit information under certain conditions to facilitate their performance, and because the results of our simple model did not conform perfectly to the human data suggests that under most conditions, the implicit and

explicit knowledge systems may interact. This includes the possibility that the structure of the implicit knowledge system may be modified by the explicit system. In terms of the model, it must be remembered that the model is restricted to coding the stimuli in isolation from any other functions that are performed by a real neural net. Therefore, it is possible that systems outside that modelled here are responsible for the effect of explicit instructions found by Reber, et al (1980). Adding such systems to the model may allow it to explain how implicit and explicit systems interact. The model may best be described as simulating both how knowledge is stored implicitly and the form of such knowledge, while also being suggestive of how implicit knowledge may be used by a verbalizable knowledge system.

Acknowledgments

We thank Prof. Marta Kutas and Prof. Mark. St John for helpful comments and discussion.

References

- Brown, T.H., Kairiss, E.W., and Keenan, C.L. 1990. Hebbian Synapses: Biophysical Mechanisms and Algorithms. *Ann. Rev. of Neurosci.* 13:475-511.
- Cleeremans, A. and McClelland, J.L. 1991. Learning the Structure of Event Sequences. *J. of Exp. Psych., Gen.* 120(3):235-253.
- Hebb, D.O. 1949. *The Organization of Behavior*. Wiley Press, New York.
- Posner, M.I. and Keele, S.W. 1968. On the Genesis of Abstract Ideas. *J. of Exp. Psych.*, 77(3):353-363.
- Reber, A.S. 1967. Implicit Learning of Artificial Grammars. *J. of Verb. Learn., Verb. Behav.*, 77:317-327.
- Reber, A.S., Kassin, S.M., Lewis, S., & Cantor, G.W. 1980. On the relationship between implicit and explicit modes in the learning of a complex rule structure. *J. of Exp. Psych., Hum. Learn., Mem.* 6:492-502.
- Reber, A.S. 1989. Implicit Learning and Tacit Knowledge. *J. of Exp. Psych.: Gen.* 118(3):219-235.
- Rumelhart, D. E. and McClelland, J. L. 1986. *Parallel distributed processing : explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.
- Servan-Schreiber, E. and Anderson, J.R. 1990. Learning Artificial Grammars With Competitive Chunking. *J. of Exp. Psych.: Learn., Mem., Cog.* 16(4):592-608.
- Squire, L.R. *Memory and Brain*, Oxford Univ. Press, New York, (1987).
- Vokey, J.R. and Brooks, L.R. 1992. The Salience of Item Knowledge in Learning Artificial Grammars. *J. of Exp. Psych.: Learn., Mem., Cog.* Forthcoming.