

An Anatomically Structured Sensory-Motor Sequence Learning System Displays Some General Linguistic Capacities

Peter F. Dominey

Vision et Motricité, INSERM Unité 94, Bron, France; and Institut des Sciences Cognitives, CNRS UPR 9075, Lyon, France

The capacity in primates to master temporal–sequential constraints from the external world might provide a basis for accommodating similar constraints in language. While the neural specialization required for language clearly distinguishes man from the other primates, it is less clear to what extent this specialization constitutes a drastic neurophysiological divergence versus a variation on an existing sequencing capability. In an effort to address this issue, an anatomically structured neural network model, previously developed to reproduce complex sensory–motor sequences and the corresponding single-unit recordings from primate prefrontal cortex, is studied in a simple linguistic context. The model is presented sentences from a small language and demonstrates a simple capacity to “understand” and generalize at different levels. Interactions between variations on (a) the model architecture and (b) the target language structure agree with data from crosslinguistic aphasia studies. These results support the hypothesis that a brain architecture for nonlinguistic cognitive functions (in this case sensory–motor sequencing) can provide a basis for a general sequence processing component of linguistic function. © 1997 Academic Press

Press

INTRODUCTION

In the ongoing controversy over the degree of general versus specialized linguistic capacities in man (e.g., Greenfield, 1991; Piattelli-Palmarini, 1994), we are fairly safe in assuming that within an inherent, nonlearned linguistic capacity (i.e., derived from brain architecture), there are components specific

This work was supported by grants to P. Dominey from the Fyssen Foundation, the France Parkinson Association, the Fondation pour la Recherche Medicale, the Direction de la Recherche et des Etudes Doctorales (DRED), the Human Frontiers Organization, and by Vision et Motricité, INSERM Unité 94, directed by Professor Marc Jeannerod. The author gratefully acknowledges the useful comments of several anonymous reviewers, Dr. Steve Small, and the Editor.

Address reprint requests to Peter F. Dominey, Vision et Motricité, INSERM Unité 94, 16 Avenue du Doyen Lépine, Bron 69500, France. E-mail: dominey@lyon151.inserm.fr.

for language processing and other generalized capacities that may apply in (and derive from) other cognitive domains (Jackendoff, 1987). Indeed, analogies have been made between cognitive elements in visually guided, goal-directed motor behavior (Arbib 1982) and the related spatiotemporal aspects of linguistic behavior.

Consider the sensory-motor sequence involved in peeling and eating a piece of fruit. There is a temporal structure imposed on this sequence by the physical constraint that one cannot (or should not!) swallow before chewing, nor peel before grasping. A brain architecture that ensures sequential state control in complex sensory-motor sequences might also contribute to an analogous capability for sequential aspects of language processing.

This is not a claim that, within the individual, sensory-motor learning provides a necessary precursor to the development of linguistic capabilities—an old issue that has been rather soundly discussed and rejected (see Piattelli-Palmarini, 1994). Instead, it is a claim that the brain architecture that has evolved to support the sequential organization of complex, goal-directed movements provides an architectural basis contributing to analogous functions (i.e., sequential organization) in language.

A neuroanatomical argument supporting this view has been suggested by Greenfield (1991), who promotes the idea that Broca's area, until about age 2, indifferently handles both language and motor control and then subsequently develops distinct circuits for these functions. Fuster (1991) comments on Greenfield, saying that Broca's area resides in a neural hierarchy for representing action, mediating the praxis of speech by controlling the oropharyngeal cortex, whereas adjacent supplementary motor area (SMA) and premotor cortex play similar roles for general skeletomuscular control. Following the maturational gradient indicated by myelogenesis (i.e., first motor, then premotor, and finally prefrontal cortex), ascending representations of "efferent copies" of cortical commands (motor or more abstract) in this corticofrontal hierarchy form progressively higher levels of representation, from muscle contraction parameters to abstract representations for goal-directed action (Fuster, 1991).

This characterization of progressively more complex frontal cortical systems would be incomplete without reference to the cortico-subcortical brain organization that involves interactions among the cortex, basal ganglia, and thalamus (Alexander et al., 1986; Dominey & Arbib, 1992). From this perspective, Lieberman (1991) argues that there are no discrete "language organs" restricted to precise cortical areas, as it is not so much damage to Broca's area, but damage in the subcortical pathways connecting Broca's area to PFC that produces permanent language deficits. Indeed, aphasia can be due entirely to these subcortical pathway lesions (Alexander et al., 1987). The cortico-basal ganglia loops serve specific functions and their damage can lead to variable patterns of manual, speech, syntax and cognitive deficits seen in Broca's aphasia. In this framework, Broca's area can be considered

part of a multifunction cortico-subcortical "organ" adapted to the regulation of sequential activity in several different domains, including precise sequential hand and orofacial motor control, and that "a later change involved using these brain mechanisms (which accessed complex motor control "rules") for the sequential rules of syntax" (Lieberman 1991).

The current study attempts to make a preliminary characterization of the contributions to general linguistic capacities that could derive from the sensory-motor architecture made up of cortex, basal ganglia, and thalamus. A neural network model (Dominey, 1993; Dominey et al., 1995) that was developed to explain cortical neuronal activity in primates trained to observe and reproduce sensory-motor sequences (Barone and Joseph 1989) is exposed to simple linguistic data (sentences) and asked to discriminate between these sentences in terms of their meaning. The model architecture is based on a predominant aspect of primate brain organization that involves interactions among cortex, basal ganglia, and thalamus (Alexander et al., 1986; Dominey & Arbib, 1992).

The rest of the paper will describe the model and will then characterize the contributions to generalized sequential aspects of language processing that derive from this system as revealed by four experiments. In summary, it will be shown that (1) the model can associate sentences in a language with their respective meaning. (2) The model can generalize associations to new sentences, while more abstract generalization requires evolution of the model. (3) The model displays inherent word order preferences that indicate a coupling between brain (model) architecture and the serial structure of language, which is discussed in the context of related crosslinguistic studies.

A SENSORY-MOTOR SEQUENCE LEARNING SYSTEM

Barone and Joseph (1989) studied prefrontal cortex (PFC) neurons in monkeys that had been trained to observe and remember a sequence of lighted push buttons and then reproduce the sequence by touching the buttons in the order they were initially presented. Individual PFC neurons encode components of the spatial locations of targets, their temporal order in the sequence, and their "touched" or "not yet touched" status. The ensemble of PFC unit activity thus encodes an ongoing representation of the internal state of the sequence input and execution that changes as each successive element of the sequence is first seen and then reproduced.

In an effort to understand how this cortical state representation could participate in the corticostriatal brain network that is known to play an important role in the regulation of sequential behavior (e.g., Dominey et al., 1995b), Dominey et al. (1995a) developed a model of sequence learning in which cortex provides a state representation capacity, and striatum an associative memory that binds internal states to the corresponding motor response.

The model (Dominey et al., 1995a) displayed in Fig. 1 is based on the

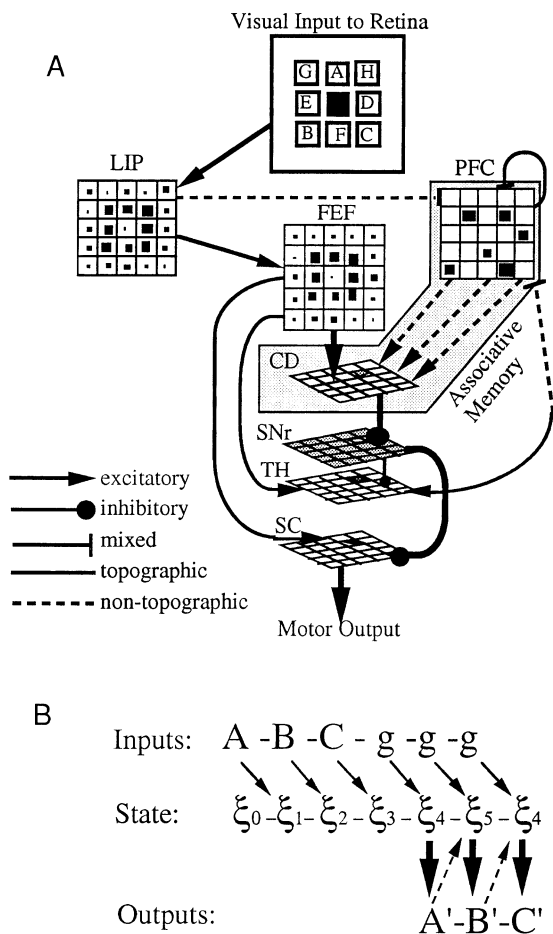


FIG. 1. Sensory-motor sequence learning model. (A) Model architecture. Each of the structures are 5×5 arrays. Sequences of spatial target locations (A-H) are presented in the (5×5) Retina input array. During output saccade generation in superior colliculus (SC), continuous inhibition of SC by substantia nigra (SNr) is temporarily interrupted by caudate (CD) inputs. CD saccade-related cells are influenced by topographic projections from frontal eye fields (FEF), and also by modifiable, non-topographic projections from prefrontal cortex (PFC). PFC combines visual input from LIP (via W^{L-P} synapses), motor efferent copy from SC (via W^{S-P} synapses), and self-input (via W^{P-P} synapses) in order to generate a time varying sequence of internal states for each presaccade period in the sequence reproduction task. These states or patterns of activity become associated with caudate activity for the correct saccade by a reinforcement learning mechanism that monitors responses and strengthens state-response connections for correct responses and weakens them for incorrect responses. (B) Schematic representation of sequence reproduction. See text.

well-known architecture of the corticostriatal oculomotor system (Dominey and Arbib 1992). Visual input to the two-dimensional (5×5 element) retina produces motor output in the form of activation of the superior colliculus (SC), a two-dimensional map of saccadic eye movements. In the simplest case, a visual input to the retina produces activation in the lateral interparietal area (LIP) of the posterior parietal cortex (Barash et al., 1991), which then activates the oculomotor frontal eye fields (FEF) (Andersen et al., 1985; Van Essen, 1979). FEF activates SC directly by an excitatory projection (Stanton et al., 1988b), and indirectly by disinhibition through the basal ganglia (Chevalier et al., 1985). SC is under tonic inhibition from the substantia nigra pars reticulata (SNr), and this inhibition can be temporarily and selectively removed when the caudate nucleus (CD) inhibits SNr based on activation of CD by FEF (Stanton et al., 1988a; Chevalier et al., 1985). Thus CD is an anatomical structure that can allow selective control of motor output choices (Rolls & Williams, 1987).

In this simplest case, the redundancy between FEF's direct and indirect (via CD and SNr) paths to the motor output in SC seems unnecessary. In the sequence reproduction task, however, at each step in the sequence, the correct one of several targets must be chosen. In this case, where a choice must be made between several targets, the brain might use PFC's selective modulation of CD to activate only the part of CD that corresponds to the correct target, ignoring the others. Indeed, PFC has the required information that uniquely encodes each successive state in the sequence (Barone and Joseph 1989) that can be used to selectively activate CD for the correct response at each successive step.

Based on evidence of learning-related plasticity in synapses connecting cortex to caudate, we implemented an associative memory in modifiable synapses between PFC and CD (Dominey et al., 1995a). The neurotransmitter dopamine (DA) is released in the striatum, following reward or events that predict reward, and participates in corticostriatal plasticity (see Schultz et al., 1995). We thus use a reinforcement learning rule so that each time a correct response is made, synapses are strengthened between the state encoding elements that are active for the current state in PFC and the units in CD that participate in the generation of the correct movement.

The anatomical basis for this associative memory exists in the projection from prefrontal cortex to striatum (Selemon & Goldman-Rakic, 1985). Kermadi et al. (1993) recorded single units in the striatum whose activity reflected this conditional association between sequence state and associated motor responses, in agreement with previous findings that striatal activity often encodes associations between internal state and behavior (Rolls & Williams, 1987).

While the details of the model can be found in Dominey et al. (1995a) and Dominey (1995), we summarize here the most important interactions. Each of the layers displayed in Fig. 1A is modeled as a 5×5 array of leaky integrator neurons. Equation (1) describes how the internal state of the

system is updated by sensory inputs from LIP, motor responses from SC (represented as routed through thalamus), and a recurrent inputs from PFC and thalamus (TH). These recurrent inputs allow the current state to influence the encoding of subsequent inputs, providing a context-dependent state transition mechanism. Equation (2) describes how the sensory-motor input to CD from FEF is biased by PFC (via adaptive W^{P-C}) to produce the learned motor response in CD. Equation (3) describes how connections between PFC and CD are modified during learning. When a response is evaluated, the connections between units encoding the current state in PFC, and units encoding the current response in CD are modified as a function of their rate of activation and learning rate R . R is positive for correct responses (Learning Rate) and negative for incorrect responses (Forgetting Rate). An effective learning strategy uses Learning Rate $>$ Forgetting Rate for initial training, and smaller Forgetting Rate $>$ Learning Rate for final training (Dominey 1995). The function $f()$ generates a non-linear output of its integrated inputs.

$$PFC = f(LIP * W^{L-P} + PFC * W^{P-P} + SC * W^{S-P} + TH) \quad (1)$$

$$CD = f(FEF + PFC * W^{P-C}) \quad (2)$$

$$W_{ij}^{P-C}(t + 1) = W_{ij}^{P-C}(t) + R * PFC_i * CD_j \quad (3)$$

Figure 1B shows the progression in time of the reproduction of sequence A-B-C. The model starts in an initial state ξ_0 . Presentation of the first visual input A drives PFC to a new state, ξ_1 . Presentation of input B drives PFC from ξ_1 to ξ_2 . Then presentation of input C drives PFC from ξ_2 to ξ_3 (Eq. 1). Presentation of the first go signal (g —illumination of A, B, and C simultaneously) drives the model to state ξ_4 and triggers the model to produce a motor output (Eq. 2) by retrieving from the associative memory the output currently associated with state ξ_4 (dark vertical arrow in Fig. 1B).

If the retrieved output is incorrect, the offending association between state encoding cells in PFC and motor response cells in CD is weakened (Eq. 3), reducing the probability that the same choice will be made again. If the output is correct, A' , this association is strengthened, and the system moves on to state ξ_5 which retrieves B' from the associative memory by the same process and so on (Eq. 3). By this trial and error learning, this system will learn the state-output associations (indicated by the thick arrows in Fig. 1B) and thus will reproduce spatiotemporal sequences as the concatenation of state-response associations.

For sequence discrimination, the model must produce a single output (rather than a sequence) in response to an input sequence. Thus, in the example above, after the presentation of an input sequence (sentence), PFC is in state ξ_4 , and the model can be trained to associate this state with the desired output (meaning). The model can thus demonstrate the capability to associate a set of sentences with their respective meanings by producing the assigned response for each sequence in response to the go signal.

In the analysis of Dominey et al. (1995a) it was demonstrated that this

model not only reproduced the behavior of sequence learning and reproduction as seen in primates, but also the single unit activity of PFC neurons during the sequence reproduction tasks. It was later demonstrated that in addition to learning simple three-element sequences, the model was also capable of learning and discriminating between complex sequences, such as *S*: ABCDABCEABCFABCG (Dominey, 1995). This sequence is complex because the repeated subsequence ABC does not have a unique successor.

Since (a) the model is built from leaky integrator neurons whose response latency is related to their input intensity, and (b) the input intensity from PFC to CD varies with learning, the model can also display learning-related reductions in response times for learned sequences. This provides the basis for simulation of serial reaction time experiments (Dominey et al. 1995c; Dominey 1997).

From a general perspective, each symbolic input to the system drives a transition to a new internal state. At the end of sequence presentation, this internal state encodes all of the content (including the temporal ordering) of the sequence of input symbols presented, similar to the “sentence gestalt” of St. John and McClelland (1990) and the context units of Elman (1990). This state information can then be used to generate outputs and/or subsequent state transitions.

The model differs from related connectionist models on several accounts: (1) Its architecture is based on that of the primate frontostriatal system which allows it to reproduce and predict neural activity. (2) Its generalized input output mapping allow it to be used for a variety of tasks including sequence learning and discrimination. (3) Learning is based on synaptic plasticity due to reinforcement signals from dopamine activity in frontostriatal connections, and occurs only in response to behavior, thus avoiding the complexity of recurrent backpropagation. (4) Activation of the CD by PFC takes place in time, depending on the strength of the associations, so reaction time is among the performance measures available for analysis.

Contributions to General Linguistic Capacity

This model can thus recognize and generate sensory motor sequences based on an ability to represent sequential context. From a linguistic perspective, this would contribute a general capacity for the recognition and generation of linguistic sequences. The analogy between the sequential structure in language and motor control predicts some specific contributions to generalized sequential aspects of language processing from a sensory–motor architecture: (1) The system should learn to associate sentences in a language with their respective meaning (based on the sequence discrimination capacity described above). (2) The system should generalize in this domain: Extend meanings of learned sentences to apply to new sentences in the same language. (3) The system’s architecture should inherently favor some syntactic structures over others. We now test these predictions.

Linguistic Experiments with a Sensory–Motor System

For the subsequent experiments a simple language that allows sentences about the transfer of possession between agents is defined and the model is exposed to sentences in that language. After presentation of such a sentence the model is trained to generate a response that indicates which agent is now in possession of the object. Various manipulations of the sentence structure etc. allow us to examine in some detail the “inherent linguistic” capacities of the model.

Syntax and semantics of a simple language. The language, **L**, uses two verbs “gives-to” and “takes-from” to specify the transfer of possession between two agents. The syntactic and semantic structures of these two verbs are described in terms of their argument structure and thematic grid. The argument structure characterizes the number and order of arguments associated with a verb, and the thematic grid gives an approximate semantic description of the sentences in which the verbs participate.

For both verbs, “gives-to” and “takes-from,” the verb is embedded between the Agent and Goal. For “gives-to,” the Agent transfers possession to the Goal, whereas the inverse is the case for “takes-from.” Possession is indicated by underlining in 1 and 2.

- | | | |
|------------------------------|----------------|---|
| 1. gives-to _v , | a-structure: | (x,y) |
| | thematic-grid: | (Agent, gives-to _v , <u>Goal</u>) |
| 2. takes-from _v , | a-structure: | (x,y) |
| | thematic-grid: | (<u>Agent</u> , takes-from _v ,Goal) |

The grammar for **L** is defined as follows:

1. $S \rightarrow x V y$
2. $V \rightarrow \text{gives-to, takes-from}$
3. $x \rightarrow A, B, C$
4. $y \rightarrow A, B, C$

The set of legal sentences in **L** is constrained such that the same element cannot occupy both the agent and goal roles. With this constraint there are 12 possible sentences in this reduced language, formed by a choice of three agents, two verbs, and two goals ($3 \times 2 \times 2 = 12$). For a given sentence “A gives-to B,” the intended meaning is that “B” now has possession of the implicit object.

Processing sentences with the model. The model allows visual input of spatial targets and produces motor outputs in the form of orienting movements to visual targets by activation of SC. For our current purposes, different visual inputs and motor outputs can be assigned to each of the linguistic symbols defined above. Figure 1 displays labeled spatial arrays of targets (A-H) that serve as the visual input to Retina, and the corresponding SC

TABLE 1
Sentences and Their “Meanings” to Be Learned
by the Model Based on Our Defined Grammar and
Semantic Structures

Sentence	Response	Sentence	Response
1. ADB	B	7. BDC	C
2. AEB	A	8. BEC	B
3. ADC	C	9. CDA	A
4. AEC	A	10. CEA	C
5. BDA	A	11. CDB	B
6. BEA	B	12. CEB	C

motor output map. The following mapping is defined between spatial targets and symbols in our language: (give-to:D, take-from:E, A-C and F-H map to their respective targets).

EXPERIMENT 1: SIMPLE LANGUAGE ACQUISITION

This experiment will test the model’s capacity to learn the associations among the 12 sentences and their respective meanings. The meaning of interest in this case is the recipient of possession. In the sentence:

(1) A gives-to B,

the recipient is B, and the model must learn to respond appropriately. Recalling our input mapping, (1) becomes:

(2) A D B,

and correct response is “B.”

Method

Using the grammar and the mapping of linguistic to sensory-motor symbols defined above for the 12 legal sentences, the model was trained to associate with each sentence its meaning—the agent that is in possession of the implicit object. Each sentence is presented as a temporal sequence of spatial input symbols as defined in Table 1. Following the presentation, the model must choose (by producing a motor output) which of the three possible agents is now in possession. As described above, the synaptic weights in W^{P-C} that associate the state in PFC with the response generation in CD are strengthened for correct responses and weakened for incorrect responses. In an initial instructed phase of learning (Pre-Train), after the input sentence, only the correct choice is provided and the model simply chooses this unique response. In the subsequent phases (Train 1–4) the model is trained to make the correct choice from among all the agents.

For example, consider sentence 1 in Table 1, A-D-B(A gave-to B), with the correct response B (B has possession). The model starts in an initial state ξ_0 . Presentation of the first visual input A drives PFC to a new state, ξ_1 . Presentation of input D drives PFC from ξ_1 to ξ_2 . Then presentation of input B drives PFC from ξ_2 to ξ_3 (Eq. 1). The model is then probed to produce the “meaning” of the sentence by presentation of the go signal (simultaneous illumination

TABLE 2
Model's Performance on Learning the Language Defined in Table 1

Phase	% correct	No. of trials	Learning rate	Forgetting rate
Pre-train	100	345	1.0×10^{-4}	2.5×10^{-5}
Train 1	86	259	5.0×10^{-6}	2.5×10^{-5}
Train 2	85	859	5.0×10^{-6}	2.5×10^{-5}
Train 3	79	850	1.0×10^{-6}	1.5×10^{-5}
Train 4	99	852	0	1.5×10^{-5}

of A, B, and C) which drives the model to state ξ_4 and triggers the model to produce a motor output (Eq. 2) by retrieving from the associative memory (in the W^{P-C} synapses) the output currently associated with state ξ_4 (dark vertical arrow in Fig. 1B).

If the retrieved output is incorrect, the offending association between state encoding cells in PFC and motor response cells in CD is weakened (Eq. 3), reducing the probability that the same choice will be made again. If the output is correct, B' , this association is strengthened. By this trial and error learning, this system will learn the state-output associations, and thus the "meanings" of the trained sentences.

This simulation is in fact a simple analog of child language acquisition in which the child (model) is told something like "John gave it to Mary" and then asked "Who has it?" If she responds correctly there is a positive reinforcement "Yes", and if not, a negative reinforcement "No."

Results and Discussion

The results are presented in Table 2 by training phase, in terms of the percent correct, number of trials, and the learning and forgetting rates (R in Eq. 3). After training, the model attained error free performance for the set of 12 sentences. Note that chance performance is 33%, since all choices are made from among A, B, and C.

The system is capable of learning a consistent and error free representation of a simple language. Clearly we are at the limit of what can be called a "language" here. However, even this simple language requires the capacity to integrate information about word order dependant on the presence of certain words (the two verbs) to encode the correct representation and response. A model that was developed to perform sequence reproduction and discrimination also displays capabilities that can be considered as contributions to generalized sequential aspects of language processing. In this case the contribution is manifest in a capacity to map symbolic strings onto their respective "meanings" in an arbitrary (but consistent) linguistic construction.

From a more general perspective the model has learned two different two-case predicates, xDy and yEx , where D (gives-to) and E (takes-from) are roughly inverses of each other in the sense that xDy has the same meaning (discounting focus etc.) as yEx . We can then ask if the model "knows" in the general case that xDy and yEx "mean" the same thing. In the set of sentences defined in Table 1, for each xDy there exists a corresponding yEx

(e.g., sentences 6 and 1). In learning all of the sentences, the model demonstrates that it “knows” that for two different two-case predicates P and P^{-1} , which share an inverse relation like “gives-to” and “takes-from” involving an agent and an effected object, xPy has the same functional meaning as $yP^{-1}x$.

EXPERIMENT 2: GENERALIZATION ON SURFACE STRUCTURE

In addition to simply learning the 12 sentences in **L**, the general linguistic capacity should also be able to generalize the meanings of the two verbs in **L** to new sentences with the same structure. Generalization normally relies on two aspects of the representation of information. First, items that are “similar” along some dimension (e.g., items of the same color) will have some corresponding degree of overlap or similarity in their representation along that dimension. That is a blue square and a blue circle will have some overlap in their internal representations; particularly along the color dimension where they share the attribute “blue.” The second requirement is that similar representations will have tendency to evoke similar responses. Once a behavioral response becomes associated with both the blue square and the blue circle, presentation of a new blue item will have a tendency to evoke this same behavior.

In the case of linguistic generalization the situation is slightly different. Consider the following templates:

(3) A E y. A. A takes-from y. The recipient is A.

(4) x D B. B. x gives-to B. The recipient is B.

One form of generalization in this domain is to learn that for all sentences of the form (3), independent of y , the recipient is A. Likewise for (4), independent of x the recipient is B. In the case of (3) (returning to generalization on color), the string “A E” is like blue, and y is like the shape, which can vary. Here, rather than considering a static spatiotemporal object like “blue square” for example, dynamic objects are constructed as a sequence of symbols. In any measure of generalization the important issue entails attending to the relevant parameter (e.g., color) while ignoring the unimportant ones (e.g. form). In this linguistic generalization task, the model must learn to ignore either the first or last element, dependant on the value of the second element.

Method

A training set of 12 sentences (G-Train) of six examples each based on the templates (3) and (4) was constructed as shown in Table 3. A testing set (G-Test), also of six examples each of (3) and (4) is defined in Table 4. In order to accommodate the number of sentences, we augmented the set of input symbols (that can occupy positions x and y in (3) and (4)) to include F-R which were assigned arbitrary unoccupied locations in the spatial array. The model

TABLE 3
Definition of Generalization Training Sentences in
G-Train

Sentence	Response	Sentence	Response
1. AEC	A	2. CDB	B
3. AEH	A	4. HDB	B
5. AEI	A	6. IDB	B
7. AEJ	A	8. JDB	B
9. AEK	A	10. KDB	B
11. AEL	A	12. LDB	B

was tested on the G-Test both in the naive state and in the trained state (i.e., after training with the G-Train). For the testing, the learning rate (R) was set to zero so that no learning would occur, and the model (naive and trained) was then tested one time each on the 12 sentences in the G-test.

Results and Discussion

The naive model performs at 42% correct on the G-Test (chance is 50% as all choices were made between A and B). After training to 100% performance on G-Train (690 trials, with Learning and Forgetting Rates 5×10^{-6} , and 5×10^{-5} respectively) the model performs at 92% correct on G-Test. This performance becomes error-free after three mistakes with the Forgetting Rate 5×10^{-5} .

The performance on the test set indicates that the system can generalize previous experience about the meaning of “give-to” and “take-from” to new sentences in the extended L. The model is capable of discriminating relevant symbols (e.g., A in (3)) in sentences from distractors (e.g., y in (3)) by a combination of their sentence position, and the “verb” that cues the relevant position. More generally, rules (3) and (4) are implicitly embedded in the sentences of G-Test and G-Train. Through exposure to the sentences the model extracts these rules, providing the basis for this generalization

TABLE 4
Definition of Generalization Testing Sentences in
G-Test

Sentence	Response	Sentence	Response
1. AEM	A	2. MDB	B
3. AEN	A	4. NDB	B
5. AEO	A	6. ODB	B
7. AEP	A	8. PDB	B
9. AEQ	A	10. QDB	B
11. AER	A	12. RDB	B

capability that would contribute to the general component of language processing (see also St. John & McClelland, 1990, and Elman, 1990).

EXPERIMENT 3: GENERALIZATION ON DEEP STRUCTURE

We now consider how the model will respond to sequences that are grammatically equivalent, but differ in their lexical elements. It has previously been described in humans, that training on letter strings generated by an artificial grammar can transfer to performance on new letter strings generated with the same grammar but different letter sets (e.g., Gomez & Schvaneveldt, 1994). In a linguistic context this corresponds to the ability to recognize that two sentences are structurally equivalent in terms of the "slots" defined by the constituent words, but that they differ in the specific word "fillers" that occupy the slots. In this experiment we examine the model's capacity to perform this kind of generalization, and consider an "evolutionary" variation on the model that aids in this function.

Because the neural elements in the model are leaky integrators with input thresholds for firing, they require time to attain an activation sufficient to produce an effective output. As specific PFC to CD synapses linking PFC states to CD motor responses get stronger through learning, the time required to drive CD and thus allow the motor output in SC is reduced, yielding a sequence specific reduction in RT. Thus, rather than using an explicit response generation, we use the serial reaction time paradigm (SRT) in order to measure the model's learning ability. In the SRT paradigm subjects perform a series of manual key presses, each in response to the successive elements of a series of visual stimuli. Nissen and Bullemer (1987) demonstrated that the reaction times for such responses are significantly reduced if the stimuli appear in a repeating sequence, as opposed to in a random order. This can be considered as a form of sequence learning because the RTs were considerably faster when the signals occurred in the trained repeating pattern rather than when they appeared randomly.

We thus use the SRT paradigm to measure the model's capacity to transfer knowledge about the syntactic structure of sentences to new sentences that share the same syntactic structure, but with different lexical items. To model an SRT task, we present a sequence of inputs to the retina, each one acting as its own "go signal," and measure the sum of the reaction times. Since there is no choice involved, there are no errors. However, after each response in the sequence, the learning rule is invoked so that that state-encoding pattern of activity in PFC at each step becomes linked with greater synaptic efficacy to the corresponding motor response in CD for that state, yielding a sequence-specific reduction in the serial reaction times (Dominey 1997).

Method

The simulated task consisted of measuring reaction times for motor responses to a series of targets appearing in one of 8 positions (A-H) on the two-dimensional spatial array Retina.

After each response the reaction time was recorded, and the appearance of the subsequent target was triggered. Blocks of trials could be of two types, sequence and random. In sequence blocks, the target locations followed a pattern of 20 elements that repeated five times for a total of 100 trials. The pattern was of the form S1: A-B-C-D-A-B-C-E-A-B-C-F-A-B-C-G-A-B-C-H. This is a complex sequence because of the ambiguity in choosing the successor to subsequence A-B-C. In the random blocks, 100 elements appeared in a random order, preserving the same element frequency as in the sequence blocks, that is, each element appeared the same number of times as in the sequence block.

In the simulation, a random ‘practice’ block was presented, followed by nine sequence blocks. Then learning was verified by presenting blocks in the order random, sequence, random, and sequence. Differences in reaction times for the random and sequence blocks were used to determine the degree of sequence learning. Finally, the degree of transfer of deep (grammatical) structure was tested by presenting a random block and then a new sequence block S2 that has the same structure (i.e., A-B-C-D-A-B-C-E-A-B-C-F-A-B-C-G-A-B-C-H) but with each letter A–H corresponding to a different spatial location than in the previous sequence. That is, the new sequence S2 is isomorphic to S1.

Results

In Fig. 2 we present the mean response times for the successive blocks described above. During the epoch labeled ‘Sequence Training’ we see the progressive reduction of reaction times. Comparing the reaction times for the sequence learning blocks to that of the first random block (R1), we can see that there is a continuous improvement (reduction) in the reaction times in successive sequence learning blocks. In the ‘Sequence Testing’ blocks this improvement can be attributed to sequence learning as it is largely eliminated when the repeated sequence S1 is replaced by a random sequence. Finally, when we look at the ‘Generalization Testing,’ we see that Sequence S2, though it has identical grammatical structure to S1, displays no benefit from training on S1 for the original model, as revealed by the reaction time near the random level.

‘*Evolution*’ of the model. In order to perform this kind of generalization, the model should recognize if elements are unique or repeats of previous elements, and learn sequence descriptions in these terms. In this sense the sequence ABABC will be the same as sequence EFEFG. Both can be described as U,U,N-2,N-2, U where U indicates a unique or nonrecognized element, and N-2 indicates repetition of the N-2nd element.

To permit this kind of recognition, the model was augmented with a set of seven short-term memory (STM) modules that invariantly remember the $n-1$, $n-2$. . . $n-7$ th previous responses (Eqs. 4.1 and 4.2) as displayed in Fig. 3.

$$STM_i = STM_{i-1} \text{ (for } i = 7 \text{ to } 2) \quad (4.1)$$

$$STM_1 = SC \quad (4.2)$$

A recognition (Recognition) function compares the current response with these STMs (Eq. 5). This recognition function is used for learning to gate

the STM contents into the output (CD) when repetition is predicted (reducing RT), and then to update the state representation in PFC (Eq. 1').

$$\text{Recognition}_i = \text{STM}_i * \text{SC} \text{ (for } i = 1 \text{ to } 7) \quad (5)$$

In order for the previous elements stored in STM to influence subsequent responses, a set of modulatory neurons (Modulation) gate the STM representations into the output, CD (Eq. 2'). PFC controls this gating in the following way. Each time a recognition occurs (Eq. 5), i.e. when the current sequence element matches a previous one, the current PFC activation becomes associated with the neuron in Modulation that gates the recognized STM element to the output, CD (Eqs. 6 and 7).

$$\text{CD} = f(\text{FEF} + \text{PFC} * \text{W}^{\text{P-C}} + \text{Modulation} * \text{STM}) \quad (2')$$

$$\text{Modulation} = \text{PFC} * \text{W}^{\text{P-M}} \quad (6)$$

$$\text{W}_{ij}^{\text{P-M}}(t + 1) = \text{W}_{ij}^{\text{P-M}}(t) + \text{R} * \text{PFC}_i * \text{Recognition}_j \quad (7)$$

$$\begin{aligned} \text{PFC} = & f(\text{LIP} * \text{W}^{\text{L-P}} + \text{PFC} * \text{W}^{\text{P-P}} \\ & + \text{SC} * \text{W}^{\text{S-P}} + \text{TH} + \text{Recognition} * \text{W}^{\text{R-P}}) \end{aligned} \quad (1')$$

Thus, PFC state is modified by the description of the previous sequence elements as unique or repeated. When a repeat is recognized, the current PFC state becomes associated with modulation of the matching STM element into CD output. The next time this sequence of unique vs. repeated elements occurs, PFC will modulate the appropriate STM element into CD before the response occurs, reducing the RT for this response.

Consider the processing of the second (underlined) B in the sequence ABABC. From Recognition, the PFC has received the input U,U,N-2 (corresponding to A B A), and is now in some state of activation. The SC output for B will match STM_2 , since the N-2nd response was also a B. By Eq. 7, the elements active for the current state representation in PFC will now become associated with Modulation_2 , the element that modulates the contents of STM_2 (element 'B') into CD. The result is an increase in the probability that the subsequence input of a subsequence of the form U,U,N-2 (e.g., ABA) will lead to activation by PFC of Modulation_2 , thus feeding the contents of STM_2 (B) into CD to yield a reduced reaction time. By this mechanism, after training on a sequence ABABC, the model will recognize (and be able to repeat and respond with reduced SRT) any sequence of the form U,U,N-2,N-2,U.

The result of this "evolution" as revealed by performance on the deep structure generalization experiment is displayed in Fig. 2 (labeled Updated Model). Note that the two model variations perform well in the sequence training and testing phases, but in the generalization testing, only the model variation that takes into account the abstract structure of sequences is able to transfer its training from S1 to S2.

Discussion

In these simulations we see a case in which the model fails to demonstrate generalization of knowledge about grammatical structure to a new set of lexical items. By introducing a mechanism that encodes sequential structure in terms of unique vs. repeated elements, i.e., in terms of a general syntactic structure as opposed to the specific surface structure, the model is able to perform this generalization to new sets of lexical items. This means that after exposure to a given sentential form, the model can recognize new sequences that have this same form. The linguistic analog is the capacity to recognize and exploit the commonality between sentences of the same grammatical form.

While other models have demonstrated related abilities to generalize the meanings (or lexical categories) of single new words based on the surrounding context (St. John and McClelland 1990, Elman 1990), they have not demonstrated a capacity, seen here, to generalize sentence forms to new sets of lexical items, an important property of human language. This is because these models can classify new words only via a context built up by familiar words, rather than by relations *between* words, familiar or new.

We can consider this an “evolution” of the model in terms of Fuster’s (1989, 1991) argument that increasingly abstract functions are implemented through the successive development of prefrontal areas that operate on the outputs from subordinate cortical areas. Specifically in this model, a new “prefrontal area” compares the motor output to a short term memory of the previous outputs to create a representation of the abstract repetitive sequential structure. This variation of the model represents an evolutionary change that accommodates an increased level of abstraction in the model’s sequence processing capacities. The motor analog to this increase in complexity is the additional capacity to perform new motor sequences with little or no additional training if they have the same structure as previously learned sequences (Dominey et al., 1995c).

EXPERIMENT 4: “INNATE” (Architecture Dependent) CAPABILITIES

Crosslinguistic studies in aphasic subjects allow the potential dissociation of performance effects due to universal (brain) linguistic mechanisms from performance effects that are language-specific (Bates and Wulfeck 1991). It is claimed here that the proposed sensory–motor network organization can contribute to a general (universal) component of linguistic capability for temporal processing. One way to examine this contribution is to perform a kind of crosslinguistic experiment by exposing variations of the model to a variety of argument structures or languages. If the native architecture has a specific contribution, it will be manifest in a preference for one or more of these structures over the others.

To address the “patients” aspect of crosslinguistic studies, the internal

TABLE 5
Sentences for Six Different Argument Structures

xDy. y	xyD. y	Dxy. y	yDx. y	yxD. y	Dyx. y
ADB. B	ABD. B	DAB. B	BDA. B	BAD. B	DBA. B
ADC. C	ACD. C	DAC. C	CDA. C	CAD. C	DCA. C
BDA. A	BAD. A	DBA. A	ADB. A	ABD. A	CAB. A
BDC. C	BCD. C	DBC. C	CDB. C	CBD. C	DCB. C
CDA. A	CAD. A	DCA. A	ADC. A	ACD. A	DCA. A
CDB. B	CBD. B	DCB. B	BDC. B	BCD. B	DBC. B

connectivity of the architecture can be modified (by re-initializing the values in the fixed connection matrices W^{L-P} , W^{P-P} , W^{S-P}), yielding a number of variations (“patients”) while leaving the gross structure intact. If there are architecture-specific capacities, they should be manifest in a general change in performance for the different argument structures resulting from these internal connectivity changes.

Method

For experiments 1 and 2, fixed argument structures were used for both verbs. For the verb “gives-to” the argument structure used was “Agent-Verb-Goal.” Here the six possible variations on this argument structure with the verb “gives-to” are tested to determine if the architecture favors one argument structure over the others:

- (1) x D y, 2) x y D, 3) D x y, 4) y D x, 5) y x D, 6) D y x.

For each of these argument structures (1–6), the six possible sentences using agents A, B and C (to replace x and y) are formed, as shown in Table 5.

At the same time we examine how variations in the model may affect argument structure preferences. To do this, the internal connectivity is modified by re-initializing the values of connection matrices W^{L-P} , W^{P-P} , W^{S-P} to produce a new variation on the architecture. A total of 10 different variations of the original model are thus formed. Then each variation is tested on the six argument structures (36 sentences) as described above to determine if variations within the overall architecture can produce modifications in argument structure preferences. During testing with a given argument structure, the model is first pre-trained on the six sentences with no choice (84 trials, Learning Rate = 7.5×10^{-4}) and then tested in the choice condition (84 trials, Forgetting Rate = 5×10^{-5}).

Results

Figure 4A displays the performance attained for each of the 10 architectures (connectivity variations) averaged over the 6 variations on argument structure. These 10 architectural variants are divided into “intact” and “impaired” groups based on their relative performance. Figure 4B illustrates the level of performance attained by these intact and impaired group architecture variations for each argument structure.

We first performed a two-way ANOVA on these data, with Architecture and Argument Structure as the two main factors. For each combination of

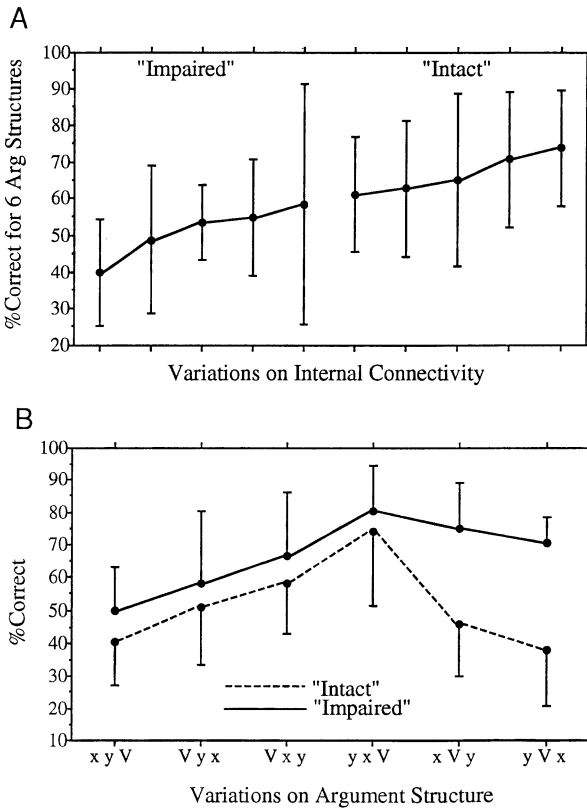


FIG. 4. Interaction between Architecture and Argument Structure. The six possible syntactic structures are tested using 10 different configurations of the fixed connections $W^{1,S}$, $W^{S,S}$, $W^{S,M}$. (A) Mean \pm SD percent correct for the 10 connectivity variations as a function of different argument structures. Intact vs. Impaired based on relative performance. (B) Mean \pm SD percent correct for the six Argument Structures as a function of connectivity for the Intact and Impaired groups.

architecture (10) and argument structure (6), the individual data for the six different sentences were used, for a total of 360 separate measures. Main effects were found for Argument Structure $F(5, 300) = 7.86$, $p < 0.0001$, and for Architecture $F(9, 300) = 4.13$, $p < 0.0001$. There was also a significant interaction between Architecture and Argument Structure, $F(45, 300) = 1.94$, $p < 0.001$. This allows us to conclude that the general architecture is sensitive to word order in argument structure, and that variations on the architecture produce selective variations in this sensitivity.

A second ANOVA was then performed with Argument Structure and Aphasia (Intact vs. Impaired from Fig. 4b) as the main factors. Again, a main effect was found for Argument Structure $F(5, 348) = 7.05$, $p < 0.0001$,

and also for Aphasia $F(1, 348) = 22.1, p < 0.0001$. There was a weak but nonsignificant interaction between Argument Structure and Aphasia $F(5, 348) = 2.16, p = 0.0585$. This allows us to conclude that, again, the architecture is sensitive to argument structure, and that the 10 variations we tested can be divided into “intact” and “impaired” performance groups which differ significantly in their performance. The nonsignificant interaction between Aphasia and Argument Structure indicates that the models’ sensitivity to word order for the different Argument structures (languages) does not significantly vary between the intact and impaired states. This will be of particular interest when considered against related cross-linguistic studies in man.

Discussion

One of the main architectural features of this model is that it is recurrent, i.e., the internal state that influences the current behavior is influenced by all previous behavior. Thus, by its construction, this architecture is sensitive to sequential order, and any performance measures reflect this sensitivity to sequential “word” order. We exposed the sensory–motor system to a variety of argument structures and considered that if the gross architecture has a specific contribution to language processing, it will be manifest in a preference for one or more of these argument structures over the others. The significant main effect for Argument Structure (see Fig. 4B) indicates indeed that this is the case. Likewise, by modifying the internal connectivity while leaving the gross structure intact, differences in architecture-specific capacities should be manifest in a change in the performance and preferences for the ensemble of Argument Structures. This was observed in the significant main effect for Architecture (see Fig. 4A). Finally, the significant interaction between Architecture and Argument Structure demonstrates that while the overall architecture is sensitive to argument structure, the specific “preferences” are modulated by changes in the initial conditions for the architecture.

If we entertain the consideration that the different argument structures correspond to different languages, and that the different connectivity variations correspond to different brain configurations due to individual differences including lesions (albeit in a limited sense), then we can attempt to interpret these results in terms of current neurolinguistic findings. In order to examine the effects of “aphasia” and argument structure and their interactions, we divided the 10 architectural variants into two “intact” and “impaired” groups based on their performance. We then determined that, by definition, the impaired and intact groups differed in their overall performance, and that both groups showed performance dependencies on argument structure. The interesting point was that the language-related (Argument Structure) differences did not depend on aphasia (Intact or Impaired) as revealed by the lack of interaction between Aphasia and Argument Structure.

This is in agreement with crosslinguistic studies in normal and aphasic subjects (see Bates & Wulfeck, 1991) for language production, comprehension, and grammaticality judgment; although there are variations among languages (Argument Structure), the relative robustness of word order for different languages is preserved in subject types (i.e., normal vs. aphasic).

In an event description task with English, Italian, and German subjects, Bates et al. (1988) demonstrated that while there are differences in importance of word order for normal subjects speaking these languages, these differences are preserved for fluent and nonfluent aphasics. Similarly, in a crosslinguistic study on sentence comprehension in these languages, Bates et al. (1987) demonstrated that though there were language-specific patterns of use of word order, this word order information showed no deterioration in any language or patient group. Indeed, many German and Italian patients overused canonical word order as compensation for reductions in their use of grammatical morphology. In comparing our data to those from human studies, clearly we cannot compare word order effects to other effects, e.g., grammatical morphology, that do not exist in the model. We can, however, say that different languages (Argument Structures) rely differently on word order, as revealed by performance differences, and that these differences are preserved in aphasias as revealed by the nonsignificant interaction between Argument Structure and Aphasia.

In summary, Experiment 4 demonstrated that this recurrent architecture has inherent preferences for some argument structures over others, and that these preferences are modulated by variations in the networks initial conditions. The "crosslinguistic" analysis confirmed that while the importance of word order may vary across languages as revealed by a significant main effect for Argument Structure, these variations are largely preserved in the aphasic conditions, as revealed by the lack of significant interaction between Argument Structure and Aphasia.

Related models. St. John and McClelland (1990) developed a model architecture specific for language processing based on an intermediate representation called the sentence gestalt (SG) that reproduced a number of linguistic capacities. The model consists of a sequential decoder that processes each constituent in turn to produce a sentence gestalt; and an output function that generates the output representation from the sentence gestalt.

Through training the model is capable of a number of linguistic tasks including: (1) Decoding the event described by a sentence by successively probing with each half of each role/filler pair. (2) Determining role assignment by probing with the filler to test if the correct role will be assigned. (3) Context dependent disambiguation, and concept instantiation by probing with role to determine if the correct filler, concept and semantic feature activation are produced. (4) Inference of nonspecified roles by probing with the role to see if the inferred filler is instantiated.

The statistical regularities of the training data are captured by the model,

and allow it to generalize. Words with common syntactic and semantic features develop similar representations in the inputs weights, allowing new words to generalize correctly based on context. Likewise, regularities in the syntactic and semantic organization of a training corpus allow the model to predict upcoming constituents in novel sentences.

Elman (1990) uses a simple recurrent back propagation network with a single hidden layer to solve linguistic problems in the framework of predicting successors in sentences. This work displays that a relatively simple recurrent network can extract a great deal of the implicit structure in sequences of words (formed by sequences of letters) that were generated by a fixed set of rules. In learning to predict the successors in large bodies of sequentially structured material, the internal representations reflect and predict that sequential structure. As in the St. John and McClelland model, words that have similar properties (as reflected in the word ordering in the sentences) produce similar activation of the hidden units, and a new word that appears in the same sequential locations as an old word will have a similar internal representation. Note that this reliance on regularities in surface structure context is not sufficient for generalization to new sets of lexical items (vs. single new words) generated with the same grammar, as in our Experiment 3, since the context must come from familiar elements. An important idea stressed in Elman's work is that the statistical regularities in a set of data (i.e., word boundaries, word successor probabilities, etc.) will be discovered by the network, and can be readily observed in the time-varying error signal (i.e., low error—predictable; high error—unpredictable). Because of its less specialized architecture, this model can be more readily applied to general sensory—motor problems. For example, Cleeremans and McClelland (1991) demonstrated that a modified version of the Elman (1990) model can be used to describe phenomena observed in sequence learning in the serial reaction time paradigm.

Jordan (1990) has used multilayer back propagation networks to study motor learning and performance. He demonstrates that constraints related to explicit task specifications, as well as intrinsic constraints that apply to a variety of tasks can be learned by optimizing a cost function that embodies these constraints. He further illustrates that the internal organization of the architecture is dependent on the nature of these constraints.

The common element to all of these sequence learning models is the use of recurrent connections that allow the internal state at step (n) to affect the subsequent internal state at step ($n + 1$). Such an architecture can support the capacity to resolve ambiguities based on sequential context and to generalize. One main difference among these models lies in their origins. The back-propagation network models do not attempt to allocate their architectural elements to brain regions, and they must answer to the standard argument that the back-propagation learning mechanism is "unbiological" because it requires the passage of information backward through synapses and

uses signals that must be precise and different for each neuron in the network (Mazzoni et al. 1991). In addition, some of these models have architectures that are tailored specifically for examination of linguistic issues with rather rigid input and output formats (e.g., St. John & McClelland, 1990), whereas others are more general (Elman, 1990) but still require a error-reduction training after each input element is processed.

Our model was developed under a number of constraints from neuroanatomy and neurophysiology of the primate corticostriatal system and uses a learning mechanism based on a physiological viable reinforcement learning scheme (see Mazzoni et al., 1991) that involves dopamine-regulated plasticity in the corticostriatal system (Schultz et al., 1995). The model (Dominey, Arbib, & Joseph, 1995) was initially developed to explain the complex pattern of spatiotemporal state encoding observed in the prefrontal cortex of behaving primates (Barone & Joseph, 1989), based on the synaptic plasticity between cortex and basal ganglia, modulated by reward information from the limbic system. It was subsequently demonstrated that the model was capable of learning to reproduce and discriminate multiple complex sequences (Dominey, 1995), and could reproduce human performance in serial reaction time tests (Dominey et al., 1995c; Dominey 1997).

This variety of performance capacities derives in part from the very unconstrained format for input—a temporal sequence of “targets” on a two-dimensional “retina,” and output—also a temporal series of “movements” described by activation of elements on a two-dimensional motor output map. Training comes from an environmental source after each response, similar to what would be seen in actual human learning.

GENERAL DISCUSSION

The purpose of this study was first to demonstrate the existence of contributions to a general linguistic capacity that can be made by an architecture for sensory—motor sequence learning and second to begin to characterize these contributions. Three forms in which these contributions might be made were identified, and then tested using an existing neural network model for sequence learning. The model (1) learned the mappings between sentences and their meanings in a small language, (2) was capable of generalizing at different levels in this language, and finally (3) displayed architecture dependent preferences for certain linguistic structures (i.e., argument structures) over others in agreement with related observations in human performance.

It is important to note that this was not an attempt to build a model of linguistic processing, but instead is a study of how an anatomically structured model that was developed to reproduce and discriminate complex spatiotemporal sequences displays general linguistic capacities. From the perspective of functional neuroanatomy, this capacity to reproduce complex sequences

which relies on state control and associative memory capabilities is directly derived from the general architecture of cortex and basal ganglia in primates (Alexander et al., 1986; Dominey & Arbib, 1992; Dominey et al., 1995a). The similar organization of multiple parallel circuits devoted to different functional modalities (e.g., eye movement, arm movement, "planning," etc.) but based on the same architecture (Alexander et al., 1986) agrees with the idea defended here that a general brain architecture (based on cortex, basal ganglia and thalamus interactions) that contributes to temporal-sequential aspects of sensory-motor function can provide a basis for similar functions in language. This is further supported by Lieberman's (1991) observation that cortico-basal ganglia loops serve specific functions, and that their damage can lead to variable patterns of manual, speech syntax, and cognitive deficits seen in Broca's aphasia and by Alexander's (1987) observations of subcortical aphasias.

This work supports the idea that a brain that is capable of complex sensorimotor sequence manipulation has some (but not necessarily all) of the required hardware for some of the sequential aspects of language processing. We examined the interactions between architecture and language processing by varying the architecture and the argument structure of the grammar, which allowed comparison with data obtained in human crosslinguistic studies. The fact that we observe significant variations in performance as a function of argument structure (language type) that are preserved between intact and aphasic subjects (patient group) agrees with the observation of Bates and her colleagues (Bates et al., 1991) that word order is an important cue, and its relative importance across languages is preserved in aphasias.

All of this is not to be confused with the idea that sensorimotor development in the individual is a prerequisite for language. Rather, it argues, as suggested by Fuster (1989, 1991), that the development of a brain architecture for sequential sensorimotor capacities in a species lays the foundation for the evolution of more complex but related linguistic functions. Simulating this "evolution" by adding a short-term memory element that modulates the state and associative memory based on abstract grammatical structure, we provided a capacity that allows transfer of grammatical learning to new sets of lexical items. It is of interest that this provides, more generally, the capacity to form and use simple analogical schemas (Gick & Holyoak, 1983) for sensory-motor behavior, allowing the system to transfer knowledge about sequential structure to new sequence instances (Dominey et al., 1995c).

How far can this line of reasoning go? Are the increasingly sophisticated requirements for complex goal directed planning enough to provide a basis for handling all of the complexities of language? Though we cannot definitively respond based on this study, the answer is likely to be no. Instead, this study is consistent with the idea that the "language organ," like the "arm movement organ" etc. takes its place in a neural hierarchy for repre-

senting action (Fuster, 1989, 1991), implemented as a specialization of the general distributed cortico-subcortical network architecture that has inherent capacities for context-dependent sequence processing.

REFERENCES

- Alexander, G. E., DeLong, M. R., & Strick, P. L. 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Ann. Rev. Neurosci.* **9**, 357–381.
- Alexander, M. P., Naeser, M. A., & Palumbo, C. L. 1987. Correlations of subcortical CT lesion sites and aphasia profiles. *Brain*, **110**, 961–991.
- Arbib, M. A. 1982. Perceptual-motor processes and the neural basis of language. In (M. A. Arbib, D. Caplan, & J. C. Marshall (Eds.), *Neural models of language processes*. New York: Academic Press. Chapter 3, pp. 531–552.
- Barash, S., Bracewell, R. M., Fogassi, L., Gnadt, J. W., & Andersen, R. A. 1991b. Saccade-related activity in the lateral intraparietal area II. Spatial properties. *Journal of Neurophysiology*, **66**, 1109–1124.
- Barone, P., & Joseph, J.-P. 1989. Prefrontal cortex and spatial sequencing in macaque monkey. *Experimental Brain Research*, **78**, 447–464.
- Bates, E. A., Friederici, A. D., & Wulfeck, B. B. 1987. Comprehension in aphasia: A crosslinguistic study. *Brain and Language*, **32**, 19–67.
- Bates, E. A., Friederici, A. D., & Wulfeck, B. B. 1988. On the preservation of word order in aphasia: Cross-linguistic evidence. *Brain and Language*, **33**, 323–364.
- Bates, E. A., & Wulfeck, B. B. 1991. Cross-linguistic research in aphasia: An overview. *Brain and Language*, **41**, 123–148.
- Chevalier, G., Vacher, S., Deniau, J. M., & Desban, M. 1985. Disinhibition as a basic process in the expression of striatal functions. I. The striato-nigral influence on the tecto-spinal/tecto-diencephalic neurons. *Brain Research*, **334**, 215–226.
- Cleermans, A., & McClelland, J. L. 1991. Learning the structure of event sequences. *Journal of Experimental Psychology: General*, **120**(3), 235–253.
- Dominey, P. F. 1993. *Models of spatial accuracy and conditional behavior in oculomotor control*. Ph.D. thesis, Department of Computer Science, University of Southern California, 201 pages.
- Dominey, P. F. 1995. Complex sensory-motor sequence learning based on recurrent state-representation and reinforcement learning. *Biological Cybernetics*, **73**, 265–274.
- Dominey, P. F. Influence of temporal organization on transfer in sequence learning: Comments on Stadler (1995) and Curran and Keele (1993). *Journal of Experimental Psychology: Learning, Memory and Cognition*, in press.
- Dominey, P. F., & Arbib, M. A. 1992. A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex*, **2**, 153–175.
- Dominey, P. F., Arbib, M. A., & Joseph, J. P. 1995a. A model of cortico-striatal plasticity for learning associations and sequences. *Journal of Cognitive Neuroscience*, **7**(3), 311–336.
- Dominey, P. F., Decety, J., Broussolle, E., Chazot, G., & Jeannerod, M. 1995b. Motor imagery of a lateralized sequential task is asymmetrically slowed in hemi-Parkinson's patients. *Neuropsychologia*, **33**(6), 727–741.
- Dominey, P. F., Ventre-Dominey, J., Broussolle, E., & Jeannerod, M. 1995c. Analogical transfer in sequence learning: Human and neural-network models of fronto-striatal function. *Annals New York Academy of Sciences*, **769**, 369–373.
- Elman, J. L. 1990. Finding structure in time. *Cognitive Science*, **14**, 179–211.
- Fuster, J. M. 1989. *The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe*. New York: Raven Press. Second edition.

- Fuster, J. M. 1991. Up and down the frontal hierarchies: Whither Broca's area? Commentary on Greenfield. *Beh. Br. Sci.* **14**, 558.
- Gick, M. L., & Holyoak, K. J. 1983. Schema induction and Analogical Transfer. *Cognitive Psychology*, **15**, 1-38.
- Goldman-Rakic, P. S. 1987. Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In *Handbook of physiology. Chap V. The nervous system.* Vol. 9, pp. 373-417.
- Gomez, R. L., & Schvaneveldt. 1994. What is learned from artificial grammars? Transfer tests of simple association. *Journal of Experimental Psychology*, **20**(2), 396-410.
- Greenfield, P. M. 1991. Language, tools and brain: The ontogeny and phylogeny of hierarchically organized sequential behavior. *Beh. Br. Sci.*, **14**, 531-595.
- Illes, J. 1988. Language production in parkinson's disease: Acoustic and linguistic considerations. *Brain and Language*, **33**, 146-160.
- Jackendoff, R. S. 1987. *Consciousness and the computational mind.* Cambridge MA London: MIT Press.
- Jordan M. I. 1990. Learning to articulate: Sequential networks and distal constraints. In M. Jeannerod (Ed.), *Attention and performance XIII.* Hillsdale, NJ: Erlbaum.
- Kermadi, I., Jurquet, Y., Arzi, M., & Joseph, J.-P. 1993. Neural activity in the caudate nucleus of monkeys during spatial sequencing. *Exp. Br. Res.* **94**, 352-356.
- Lieberman, P. 1991. Speech and brain evolution. Commentary on Greenfield. *Beh. Br. Sci.* **14**, 566-568.
- Mazzoni, P., Andersen, R. A., & Jordan, M. I. 1991. A more biologically plausible learning rule for neural networks. *Proc. Natl. Acad. Sci. USA*, **88**, 4433-4437.
- Nissen, M. J., Bullemer, P. 1987. Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, **19**, 1-32.
- Piattelli-Palmarini, M. 1994. Ever since language and learning: afterthoughts on the Piaget-Chomsky debate. *Cognition*, **50**, 315-346.
- Rolls, E. T., Williams, G. V. 1987. Sensory and movement-related neuronal activity in different regions of the primate striatum. In Schneider & Lidsky (Eds.), *Basal Ganglia and motor behavior.* pp. 37-59.
- Schultz, W. A., Romo, R., Ljungberg, T., Mirenowicz, J., Hollerman, J. R., & Dickinson, A. 1995. Reward-related signals carried by dopamine neurons. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia,* MIT Cambridge, MA: MIT Press. Pp. 223-243.
- Selemon, L. D., & Goldman-Rakic, P. S. 1985. Longitudinal topography and interdigitation of corticostriatal projections in the Rhesus monkey. *Journal of Neuroscience*, **5**, 776-794.
- Shapiro, L. P., McNamara, P., Zurif, E., Lanzone, S., & Cermak, L. 1992. Processing complexity and sentence memory: Evidence from amnesia. *Brain and Language*, **42**, 431-453.
- St John, M. F., & McClelland, J. L. 1990. Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, **46**, 217-257.
- Stanton, G. B., Goldberg, M. E., & Bruce, C. J. 1988a. Frontal eye field efferents in the Macaque monkey: I. Subcortical pathways and topography of striatal and thalamic terminal fields. *J. Comp. Neurol.* **271**, 473-492.
- Stanton, G. B., Goldberg, M. E., & Bruce, C. J. 1988b. Frontal eye field efferents in the Macaque monkey: II. Topography of terminal fields in midbrain and pons. *J. Comp. Neurol.* **271**, 493-506.
- Van Essen, D. C. 1979. Visual areas of the mammalian cerebral cortex. *Annu. Rev. Neurosci.* **2**, 227-263.