

C. De Lillo · D. Floreano · F. Antinucci

Transitive choices by a simple, fully connected, backpropagation neural network: implications for the comparative study of transitive inference.

Received: 7 August 2001 / Accepted after revision: 29 May 2001 / Published online: 29 June 2001

© Springer-Verlag 2001

Abstract In search of the minimal requirements for transitive reasoning, a simple neural network was trained and tested on the non-verbal version of the conventional “five-term-series task” – a paradigm used with human adults, children and a variety of non-human species. The transitive performance of the network was analogous in several aspects to that reported for children and animals. The three effects usually associated with transitive choices i.e. “symbolic distance”, “lexical marking” and “end-anchor”, were also clearly shown by the neural network. In a second experiment, where the training conditions were manipulated, the network failed to match the behavioural pattern reported for human adults in the test following an ordered presentation of the premises. However, it mimicked young children’s performance when tested with a novel comparison term. Although we do not intend to suggest a new model of transitive inference, we conclude, in line with other authors, that a simple error-correcting rule can generate transitive behaviour similar to the choice pattern of children and animals in the binary form of the five-term-series task without requiring high-order logical or paralogical abilities. The analysis of the training history and of the final internal structure of the network reveals the associative strategy employed. However, our results indicate that the scope of the associative strategy used by the network might be limited. The extent to which the conventional five-term-series task, in absence of appropriate manipulations of training and testing conditions,

is suitable to detect cognitive differences across species is also discussed on the basis of our results.

Key words Transitive inference · Neural networks · Five-term-series task · Cognition

Introduction

The implications of the finding that very young children (Bryant and Trabasso 1977; de Boysson-Bardies and O’Regan 1973; Riley and Trabasso 1974), apes (Gillan 1981), monkeys (McGonigle and Chalmers 1977, 1992), pigeons (von Fersen et al. 1991), rats (Davis 1992; Roberts and Phelps 1994; Dusek and Eichenbaum 1997) and hooded crows (Lazareva et al. 2000) respond transitively in the five-term-series task are still controversially debated.

As originally devised by Bryant and Trabasso (1971), the five-term-series task involved a training phase where five rods, differing in length and colour, were presented in pairs (the premises $A > B$, $B > C$, $C > D$, $D > E$) to young children. The rods protruded from a box, so that their actual length could not be perceived and the children were required to use the colour differences to choose (pointing to the selected item) between different lengths. Following a choice, either the length of the rods was shown or a verbal statement was given as feedback. After the training, the subjects were tested on all the possible pair combinations of the five terms without receiving any differential feedback after their choice. Children as young as four showed a transitive bias in their choices in the critical comparison $B > D$ (the only non-adjacent items equally referred to as “longer” and “shorter” during training).

Successively, starting with the study by McGonigle and Chalmers (1977), non-verbal versions of the paradigm have been devised for animal studies. Although the paradigm has been implemented using different stimuli and apparati in order to meet the ergonomic requirements of the species under scrutiny, the non-verbal versions of the five-term-series task share the same general proce-

C. De Lillo (✉)
Department of Psychology, University of Leicester,
University Road, Leicester LE1 7RH, UK
e-mail: cdl2@le.ac.uk,
Tel.: +44-116-2525155, Fax: +44-116-2522705

D. Floreano
Evolutionary and Adaptive Systems, Institute of Robotics,
Swiss Federal Institute of Technology (EPFL),
1015 Lausanne, Switzerland

F. Antinucci
Istituto di Psicologia C.N.R., Via Aldrovandi 16b,
00197 Roma, Italy

ture. Typically, the subjects are initially trained to criterion on a set of four binary discriminations, $A+B-$, $B+C-$, $C+D-$, and $D+E-$, where plus and minus indicate the rewarded and the unrewarded stimulus, respectively. Then, they are tested (without differential reinforcement) on novel pairs obtained from all the possible permutations of the set ABCDE. The strongest evidence for transitive inference is shown again by the choice of term B in the pair B-D, whose terms are non-adjacent, equally reinforced and non-reinforced during training, but are at different ordinal positions in the series $A>B>C>D>E$.

In addition to transitive choices, three main effects, emerging in reaction time analyses (Trabasso et al. 1975; McGonigle and Chalmers 1992) or in the percentages of correct choices (von Fersen et al. 1991), have been consistently reported in the literature. The first is the “symbolic distance effect” (these labels reflect interpretations of the effects in the early literature and here we have decided to keep them without committing ourselves to those interpretations): the greater the distance between the elements of the pair compared, in terms of number of items separating them according to the sequence $A>B>C>D>E$, the faster/better the comparison. The second has been dubbed the “lexical marking effect”: comparisons involving the first term A are responded to faster/better than the comparison involving the last term E. The third is the “end-anchor effect”: comparisons involving one of the end terms (A or E) are responded to faster/better than comparisons involving intermediate terms.

Early results obtained with very young children have been taken as a serious challenge to Piaget’s theory (see Breslow 1981). According to Piaget (1928) transitivity is the final result of a series of steps in the development of reasoning and involves logical operations. Since very young children show transitive behaviour in the five-term-series task, some authors have even suggested that logic could be innate (see Breslow 1981 for a review of early interpretations of the finding). The paradigm has been extended to non-human species in an attempt to provide more parsimonious explanations of transitive inference (McGonigle and Chalmers 1977) and none of the species tested to date failed to show evidence for transitivity.

Various authors have attempted to explain the performance of children and animals in terms of choice models that do not require logical abilities (de Boysson-Bardies and O’Regan 1973; McGonigle and Chalmers 1977, 1992; von Fersen et al. 1991; Harris and McGonigle 1994). A set of simple conditioning rules has also been proposed to account for the development of transitive choices in the B-D pair, based on the observation that in some experiments the animals experienced different reinforcement histories for each of the premises (Couvillon and Bitterman 1992; Wynne et al. 1992). However, the extent to which reinforcement history accounts of transitive performance are suitable to explain the experimental data under a variety of training and testing conditions is still not clear (see Wynne 1997; Steirn et al. 1995; Clement and Zentall 2000, for different views on this mat-

ter) even when data concerning pigeons alone are examined.

When the results obtained with a broader range of species are considered, the variety of possible explanations of transitive performance increases and several issues still have to be clarified. For example, does the training induce an ordinal representation of the stimuli in some of the species studied? Can similar learning mechanisms account for the performance of birds, rats, non-human primates, young children and human adults alike?

In serial learning and transitive inference studies, convincing arguments have been presented which support a distinction between the learning processes and the form of representation underpinning transitive choices in primate and avian species (Terrace 1993; Terrace and McGonigle 1994; McGonigle and Chalmers 1992; De Lillo 1996). Nevertheless, other authors suggest that very similar competencies might underpin transitive performance in all the different animal species tested so far and possibly in children and adult humans (Delius and Siemann 1998; Wynne 1995).

This study is an attempt to provide a new heuristic tool with which such questions can be addressed and which can be used to inform future empirical studies on the most (or least) promising paradigms for the comparative study of transitive inference. We devised a simple neural network which was trained and tested using experimental procedures which were made as similar as possible to the non-verbal version of the task presented to humans and animals. Our aim was to assess the most elementary conditions for successful transitivity using an “organism” equipped with minimal design features and competencies. We reasoned that an analysis of the relative robustness of the performance of such a device under variations of the training and testing conditions would provide useful information regarding the manipulations which are most likely to enable the detection of qualitative interspecies or developmental differences. Compared with other simulation approaches – such as those described by Couvillon and Bitterman (1992) and by Wynne et al. (1992) – neural networks provide the advantage of allowing us to evaluate the internal representation of the task. Moreover, the dynamic properties of neural networks make it possible to observe the development of this representation during task exposure as well as the outcome of different training histories.

The use of neural networks in the study of transitive inference has precedents in modelling work carried out by Delius and Siemann (1998; Siemann and Delius 1998). Aiming at producing a realistic model of transitive inference in biological systems, these Authors implemented Luce’s learning model (Luce 1959) as a relatively complex neural network with specific excitatory and inhibitory connections and weight distribution designed beforehand to produce the required outcome.

By contrast, our goal was not to develop a distributed (PDP) model of transitive inference, but rather, by using a fairly conventional, fully connected, backpropagation network, to reduce to a minimum the constraints imposed by the programmer on the structure of the network a priori.

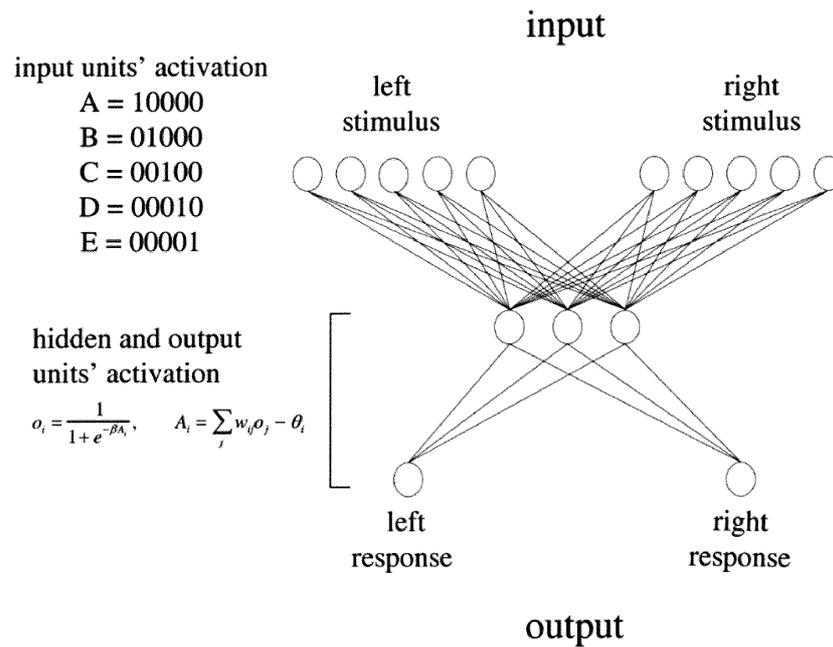


Fig.1 The architecture of the neural network: input units (top layer), output units (bottom layer) and hidden units (middle layer). The ten input units can be considered as two sets of five units, each set identifying one of the two stimuli presented simultaneously on the left (left stimulus) and on the right (right stimulus). The five digit strings used to activate the input units according to the particular stimulus presented are shown at the top left of the figure. The activation of the two output units was considered as the response of the network to the stimulus on the left (left response) and on the right (right response), respectively. The formula describes the sigmoid transfer function used to determine the activation of the hidden and output units on the basis of incoming activation from the preceding unit layer (j). $w_{(ij)}$ are the connection strengths (in the range ± 8), β is a constant (here held to 1), and θ_i is the adjustable threshold of the node. The output was thus a continuous value within the range [0, 1]

By doing this we aimed to detect the minimal requirements of transitive inference tasks, to identify the simplest strategies for their solution and to evaluate the generality of these accounts by manipulating training histories and post-training tests.

General method

A feed-forward network featuring ten input units, three hidden units, and two output units was employed. Each input unit was connected with all hidden units, and each of the hidden units was connected with the two output units (Fig. 1). The activation o_i of node i (except for input nodes whose activity was determined by the stimulus presented) was given by the weighted sum of its inputs filtered through a sigmoid transfer function:

$$o_i = \frac{1}{1 + e^{-\beta A_i}}, \quad A_i = \sum_j w_{ij} o_j - \theta_i$$

where j is the index of the nodes in the layer below, w_{ij} are the connection strengths (in the range ± 8), β is a con-

stant (here held to 1), and θ_i is the adjustable threshold of the node. The output was thus a continuous value within the range [0, 1].

This particular architecture was selected in order to make the presentation of the stimuli and the recording of the network's responses as similar as possible to the procedure generally reported in the experimental literature. Besides its architecture, we did not want to provide the network with built-in additional knowledge by constraining the system in any way. Hence, ten input units were necessary for presenting each of the premises (the input layer can be considered as two sets of five units, each set identifying the presence of one of the two stimuli of a pair) and two output units were the minimum required to indicate the selected term. From a number of parametric studies, we determined that the number of internal units was irrelevant for the main issues addressed here, affecting only slightly the speed of learning. Finally, the choice of full connectivity between layers was made in order to avoid additional externally imposed constraints on the architecture.

The stimuli, derived from a five-term set, were presented as five-bit strings of four 0s and a 1. The position of the 1 in the string indicated a particular term (A=10000; B=01000; C=00100; D=00010; and E=00001). This form of stimulus presentation ensured that initially all the terms had an equal influence on the activation of the internal units and therefore was adopted in order to preserve the arbitrary relationship between the stimuli and their ordinal position in the series A>B>C>D>E as prescribed by the paradigm. Both members of a pair were presented simultaneously, one to the left set of input units (Fig. 1, left stimulus) and one to the right set of input units (Fig. 1, right stimulus).

The two output units indicated the response for a given stimulus pair, by "pointing" to the position of one of the two input sets; for example, the correct response for the

pair A+ (left) and B- (right) would be an activation of 1 for the output unit on the left (Fig. 1, left response) and an activation of 0 for the output unit on the right (Fig. 1, right response), whereas for the pair B- (left stimulus) and A+ (right stimulus) would be 0 (left response) and 1 (right response). Having the network's response indicated by the activation of two output units (acting as "pointers" to the position of the selected item) was considered analogous to the choices of an experimental subjects in non-verbal versions of the task (where children had to point to a rod, monkeys had to reach for an object, and pigeons had to peck a key), in contrast with the alternative of having the network reproducing the selected item (coded by five output units), which could be interpreted as analogous to its verbal description irrespective of position.

In analogy with the model of Couvillon and Bitterman (1992), the connection weights were changed according to the method of back-propagation of error (Rumelhart et al. 1986), which is an extension of the correction rule of Rescorla and Wagner (1972). In both cases, the parameters leading to a given response are modified proportionally to the discrepancy between the desired and the observed response. The back-propagation procedure (Rumelhart et al. 1986) employed here is a derived general technique which modifies the parameters (connection weights) of multi-level and continuous non-linear processing systems. The algorithm works as follows. An input pair is presented to the network and the unit activations are computed up to the output layer; the response of the network is compared with the correct response and an error measure, the pattern sum of squares error (PSSerror) is generated for the output units:

$$\text{PSSerror} = \sum_i (z_i^\mu - o_i^\mu)^2$$

where z_i is the desired response and o_i the observed response of node i for the input pair μ . So, for example, in a particular trial where the pair A+ (left) B- (right) has been presented, the desired output would be left response (LR)=1, and right response (RR)=0. If the observed output is LR=0.7, and RR=0.3, then the PSS error would have a value of 0.18 i.e. $(0.7-1)^2 + (0.3-0)^2$. This error measure is propagated backwards through the connection weights in order to calculate the contribution to the error by each internal node; all the connection strengths are then modified according to the back-propagation rule:

$$w_{ij}^\mu = w_{ij}^{\mu-1} + \eta \Delta w_{ij}^\mu + \alpha \Delta w_{ij}^{\mu-1}$$

where Δw_{ij} is the weight change, η is the learning rate and α is a parameter which controls the fraction of the previous weight change (it adds a sort of inertia in the search and is, hence, also called "momentum"). Parametrical studies showed that the main findings are not altered by a particular choice of these two constants (as long as $0 < \eta, \alpha \leq 1$) which mainly affect the speed of learning.

Experiment 1

In this experiment we implemented the training procedure of the conventional five-term-series task with five neural networks (or five simulation runs in a different terminology), differing solely in the initial random connection strengths. The experiment aimed to assess the ability of the networks to perform transitive choices, and to display the main effects that are typically associated with transitivity in the literature.

Methods

The synaptic weight values of the five neural networks were initially set to random values in the range ± 1.0 , a condition that provided undifferentiated output for all input pairs. The training phase featured the presentation of the four pairs A+B-, B+C-, C+D-, D+E-, where the position (left or right) of each stimulus and the order of presentation of the pairs were randomised. Synaptic weights were updated after each presentation of a pair by comparing the network response with the correct response. The training phase continued until the average PSS error on the output units was reduced to 3.0×10^{-4} , corresponding to c. 1250 presentations of each pair (approximately the same number used in the simulations performed by Couvillon and Bitterman 1992). Each network differed in the initial set of random synaptic weight values and received a different random presentation of the pairs. All the other network parameters were identical. The learning rate was set to 0.7 and momentum to 0.1. After training, the connection strengths were recorded and the networks were tested on all the possible binary combinations of the items. During testing the connection strengths were not modified in analogy with the experimental literature where no differential feedback is provided in this phase.

Results

All the networks learned to respond correctly to the training pairs and showed a transitive bias in the test pair B-D. Since outputs of artificial neural networks approximate biological reality only to a limited degree and might be considered as expression of different behavioural variables, we used the PSS error as a general measure of performance and compared it with both the measures used in experiments with biological organisms, namely percentages of errors and latencies. Higher PSS errors imply a less accurate response and can be conceived as a higher probability of responding incorrectly or with a longer latency.¹

¹ Nevertheless, a direct transformation of PSS error values into percentages of correct responses (e.g. a PSS error of 0.18 directly interpreted as 82% correct) would be misleading and arbitrary. A specific interpretation in terms of latencies would be even more tentative. PSS error in the present context can be meaningfully used only to evaluate the relative performance of the networks with different stimulus pairs.

Table 1 Averaged PSS error ($\times 10^{-4}$) generated by the networks for each pairing of the stimuli A, B, C, D, E. Above: values obtained in the condition featuring a random order presentation of the stimulus pairs; below: values obtained in the condition featuring a serial order of presentation of the stimuli are reported (the pattern sum of squares error, PSS error, for undifferentiated output is 5000×10^{-4})

Term	B	C	D	E
Random order				
A	1.85	0.39	0.66	0.23
B		2.62	2.96	1.38
C			5.63	1.12
D				1.90
Serial order				
A	1.86	0.44	0.66	0.24
B		2.25	2.65	1.60
C			5.92	1.46
D				1.88

As shown in Table 1, the average performance for the pair B-D (2.96×10^{-4}) was comparable to the mean performance on the training pairs (3.0×10^{-4}); the pair C-D produced the worst performance, followed by the pair B-C, while the pairs A-B and D-E produced the highest levels of performance. The mean performance for non-adjacent testing pairs (1.12×10^{-4}) was higher than for the training pairs.

The symbolic distance effect, the lexical marking effect, and the end-anchor effect (Fig. 2a–c) were clearly displayed by the networks.

Experiment 2

De Boysson-Bardies and O'Regan (1973) compared children and adults in two conditions of training, one involving a serial presentation of the premises (i.e. A+B– first, then B+C–, and so on), and a second involving their presentation in random order. The underlying assumption is that subjects who attempt to construct a linear representation of the series should find the serial presentation of the premises easier. Their results demonstrated that adults' performance worsens in the random condition whereas children's performance is indistinguishable in the two conditions. They, therefore, proposed that children relied on a "verbal strategy", labelling the terms A and B as "big" and the terms D and E as "small", while the term C is left unlabelled and is responded to by default. This latter hypothesis was then tested with children by means of separate comparisons of a novel term (X) with each of the items A, B, C, D, and E. In accordance with the labelling hypothesis, it was found that the novel term X was selected above chance when presented with terms A and B and avoided when presented with terms D and E, while random choices were observed when the novel term was paired with the term C.

In this second experiment we implemented the procedure of de Boysson-Bardies and O'Regan (1973) as an initial attempt to evaluate how the performance of our net-

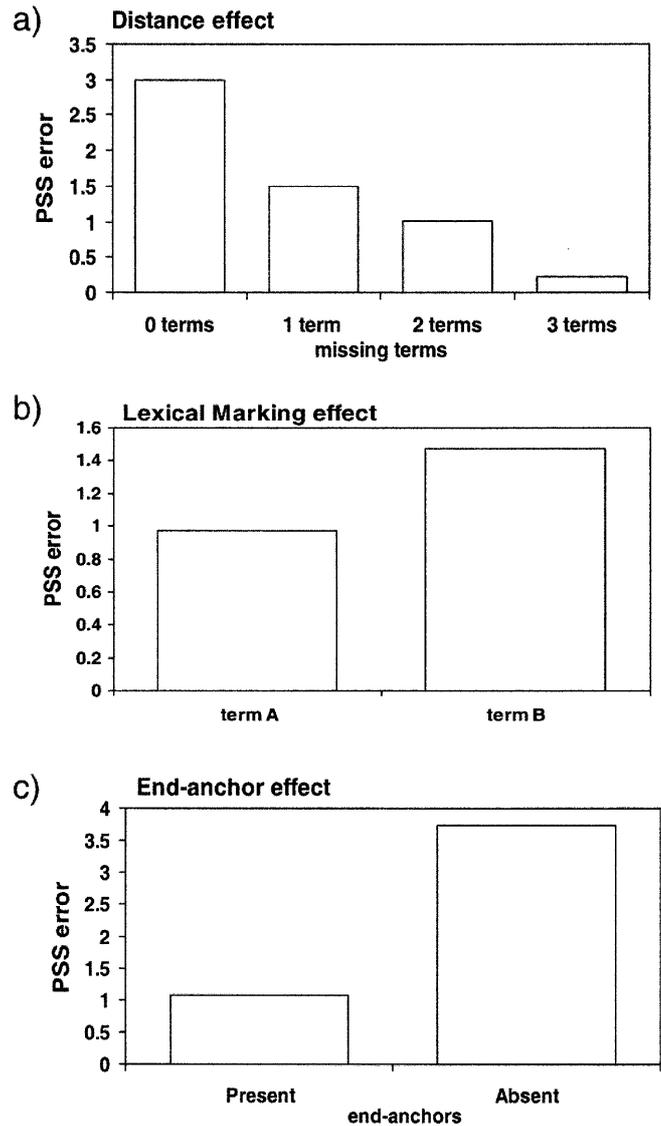


Fig. 2a–c The three effects expressed in terms of the average PSS error ($\times 10^{-4}$) for the five neural networks (simulation runs). a Symbolic distance: the pattern sum of squares error (PSS errors) were separately averaged over all the stimulus pairs with no missing terms (training pairs), with one missing term, with two missing terms, and with three missing terms. The smaller the distance, in missing terms, along ordered series A>B>C>D>E, the larger the observed PSS error. b Lexical marking: PSS errors were separately averaged over all the pairs which included the term A and over all the pairs which included the term E. Pairs including the term E generated a larger PSS error. c End anchor: PSS error separately averaged for all the pairs which included either the term A or the term E and for all the remaining pairs. A smaller PSS error can be observed for pairs which include at least one of the end-anchors

works was affected by manipulations of training and testing conditions.

Methods

The networks, starting with the same initial connection strengths employed in experiment 1, were trained with a

serial presentation of the input pairs (i.e. first A+B-, then B+C-, and so on), for the same number of times (1250).

Following training, the presentation of a new term X was implemented by introducing a new pattern of five 0s and presenting it along with other terms during the test. Thus, the new term was equally dissimilar from all the other terms by a single bit.

Results

Table 1 shows the average performance of the five networks when trained with a serial presentation of the premises. No significant difference was found (χ^2 test of the generalised median=0.4, n.s.) between this training condition and the random presentation of the training pairs (Table 1).

When the novel term X was presented after training, it was rejected (activation of the corresponding output unit close to 0) when paired with terms A and B (whose corresponding activation of the output unit was close to 1) and selected when paired with D and E, whereas, the comparison with the term C fell at chance level (the activations of both units were approximately equal). The same pattern of results was obtained when the novel term X was composed of five 1s (which is also equally distant from all the other terms, differing for 4 bits from each term).

Analysis of the microdevelopment of performance

The analysis of the development of the performance described in experiment 1 during training shows an analogy with a phenomenon found by Trabasso et al. (1975) in children: the neural networks first learned the premises which included the end-anchor terms A and E (Fig. 3). The explanation lies in the fact that these two terms are easier to learn because they are the only ones which are

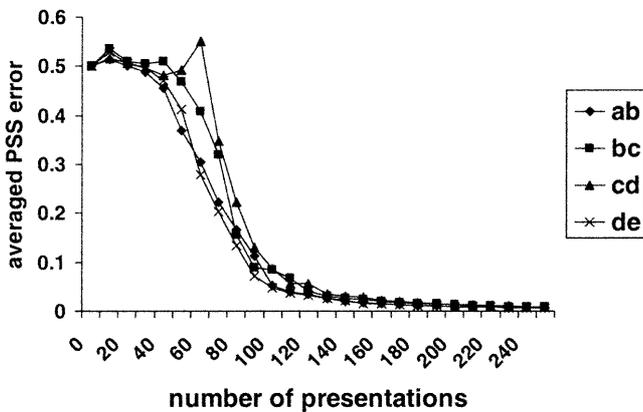


Fig. 3 Initial learning curve corresponding to the first 250 presentations, in random order, of each stimulus pair AB, BC, CD, DE. Identical learning curves are displayed by all the networks (simulation runs). Data points refer to the PSS error generated by the presentation of each of the premises during training

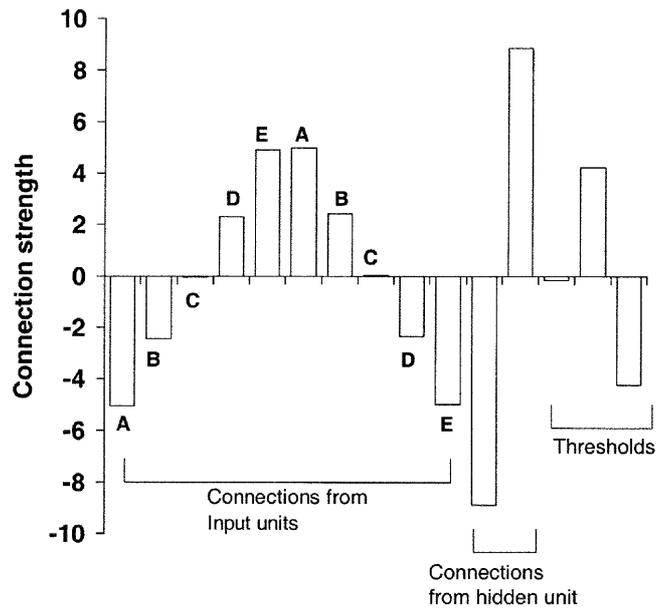


Fig. 4 Connection strengths of all the connections of a single internal (hidden) node. The first ten bars are incoming connections from input units and are named according to the stimulus identified by the position of the single one in the five-bit string presented. The following two bars come from the hidden unit to the output units. The last three bars represent the threshold values of the hidden unit and the output units, respectively. An identical structure is found for all the other hidden units of the network

always winning (A) or losing (E) (see also Wynne et al. 1992 for a similar conclusion).

An inspection of the connection strengths after training as described in experiment 1 shows how the neural networks solve the task. Figure 4 plots the strengths of the connections from the ten input units to a single internal unit and from the same unit to the output units (including the threshold values); the same configuration is repeated for all the internal units in the network (it is also formed when the networks are originally trained with a different number of internal units), irrespective of their number (in fact, neural networks with a range of one to four hidden units were analysed). This structure is formed in the early stages of training (about 250 presentations of each training pair with learning rate=0.7 and momentum=0.1 are sufficient) and becomes further consolidated with additional exposure to the stimuli.

During training the networks assign an associative strength to each term by incrementing the magnitude of the connection from the relevant bit, producing a global representation with a perfect symmetry. The terms A and E are associated with strong connection weights. The same structure develops for the terms B and D, although the absolute synaptic strength is weaker. The term C is associated with near-to-0 strength. The more distant the terms are in the series, the stronger the weighted sum that activates the hidden node becomes (the sigmoid activation function employed here ensures that strong inhibitory signals drive the hidden unit activation close to 0, whereas strong excitatory signals drive it close to 1). As a conse-

quence, the output is more differentiated with end-anchor terms; by contrast, when two terms are closer, the total amount of signal is weaker and the output is less differentiated. The choice of the output unit to be activated is then mediated by the connections from the internal node to the output nodes.

General discussion

The results of this study show that a very simple fully-connected backpropagation neural network can mimic several experimental results. The associative strength attained by each term during training was sufficient to produce transitive responding on the critical test B-D and on all the remaining binary combinations of the five terms. Although a visual inspection of the weights might suggest linear ordering, the synergistic functioning of the nodes in the network, both during learning and testing, excludes a seriation ability. Since the output is merely generated by the summation of incoming activations, the network could not establish the ordinal position of a specific stimulus in relation to others. Rather, the functioning of the network seems to offer an analogy with the labelling strategy advanced by de Boysson-Bardies and O'Regan (1973), but without resorting to high-level symbolic processes such as verbal categorical coding. The end-anchor terms A and E were initially (see the learning dynamics) "labelled" as "strong" and "weak". The rest of the associations were the built around them. The term C acquired an associative strength equal to the mean strength of all the other terms (which was around 0.0). In other words, it remained "unlabelled". This could be inferred also by the fact that the novel term X was responded to as "weaker" than A and B, and "stronger" than D and E, but an indeterminate response (both output units having similar activation) was observed in the comparison with the term C.

Moreover, as reported in the experiment of de Boysson-Bardies and O'Regan (1973) with children, no significant difference was found between the performance of the networks following a serial or a random presentation of the premises; this contrasts with results obtained with human adults, that learned better and faster when the training phase featured the presentation of the premises in serial order. The network was not pre-programmed to organise the items according to a seriation principle and our results show that a serial presentation of the premises did not lead to a better performance per se, in absence of a spontaneous (i.e. not induced by training) tendency to organise the training pairs in a principled way.

This insensitivity of the network to the order of presentation of the premises seems to highlight the schism often emerging from the experimental literature between solutions of the five-term-series, and related serial learning tasks, based on purely associative principles, as often attributed to pigeons (Terrace 1993; Terrace and McGonigle 1994; but see also Wynne 1997 for an interesting exception) and solutions based on an ordered representation of the series which most likely apply to human subjects (de

Boysson-Bardies and O'Regan 1973; Kallio 1982), non-human primates (Terrace 1993; Terrace and McGonigle 1994; D'Amato and Colombo 1988), and possibly rats under specific testing conditions (Roberts and Phelps 1994). We are currently extending the present study to investigate the extent to which the manipulation of learning parameters at different points of training might affect the sensitivity of our networks to the order of presentation of the premises.

A non cognitive account of pigeons' transitive performance in the five-term-series task invoked a positive transfer of value from the reinforced term (e.g. A+) to the non reinforced term (e.g. B-) presented as a premise to explain the preference of B over D in the critical comparison BD (von Fersen et al. 1991). Our network was not specifically programmed to allow such a value transfer. We used a simple backpropagation learning algorithm, which although regarded by some as neurally implausible (Delius and Siemann 1998), nevertheless allows us to rule out the necessity to postulate the participation of value transfer processes for the successful solution of the task. The extent to which the network can cope with other experimental procedures specifically designed to test for value transfer theory in pigeons (Steirn et al. 1995; Weaver et al. 1997) is also currently under evaluation.

Overall, our results seem to support the hypothesis formulated by Couvillon and Bitterman (1992) on the basis of their mathematical model. Simple associative rules are sufficient for success in the five-term-series task. However, the generality of this type of learning mechanism must be carefully evaluated. In our case, the network failed to reproduce the behaviour of human adults when the training condition was appropriately manipulated. Moreover, instances of dissociations between premise retention and transitive performance as observed in some studies in humans (Siemann and Delius 1996) and rats with parahippocampal lesions (Dusek and Eichenbaum 1997) must also be accounted for by theories postulating that transitive responding invariably emerges from the development of associative strength induced by reinforcement during premise acquisition.

In our opinion, however, a major limitation of the approach we adopted here is that the range of testing conditions that can be used with the network is limited to binary versions of the task, and to our knowledge no successful attempts have been reported of connectionist implementations of the task robust enough to deal with both binary and triadic versions of the task as implemented with primates (McGonigle and Chalmers 1977, 1992; see also Harris and McGonigle 1994 for a production-rule-based model that takes into account performance on triadic versions of the task). Indeed, a first attempt on our part to do so has been unsuccessful. We were unable to simulate the successful transitive performance shown by monkeys in the course of triadic testing (McGonigle and Chalmers 1977, 1992). A major challenge faced by future developments of the work presented here would therefore be a successful PDP implementation of both the binary and the triadic versions of the five-term-series problem.

In conclusion, here we have strong evidence that the output of a system that did not experience either cognitive growth or the pressure of natural selection is very similar to that of young children and animals in the five-term-series task. In fact, with a simple distributed device, such as the one we employed (apart from its basic architecture the network can be considered as a tabula rasa because of the initial random distribution of its synaptic weights), we can be sure of the absence of unsuspected built-in capacities; a possibility that, in principle, can never be excluded with experimental subjects, however young or taxonomically distant from humans they might be.

This suggests that the binary, non-verbal, five-term-series task might not be suitable for detecting ontogenetic or phylogenetic trends in the development of the cognitive skills underlying inferential abilities. In order to find behavioural differences of potential comparative significance, it might prove a more fruitful exercise to manipulate the training procedures (see also Higa and Staddon 1993; Treichler and Van Tilburg 1999) and the structure of the task itself (such as the triadic testing introduced by McGonigle and Chalmers 1977, 1992), instead of extending the same binary version of the paradigm to yet more non-human species or younger children.

Acknowledgements D. Floreano was supported by the Swiss National Science Foundation, grant no. 620-58049.99.

References

- Boysson-Bardies B de, O'Regan K (1973) What children do in spite of adults' hypotheses. *Nature* 246:531–534
- Breslow L (1981) Re-evaluation of the literature on the development of transitive inferences. *Psychol Bull* 89:325–351
- Bryant PE, Trabasso T (1971) Transitive inferences and memory in young children. *Nature* 232:456–458
- Clement TS, Zentall TR (2000) Determinants of value transfer and contrast in simultaneous discriminations by pigeons. *Anim Learn Behav* 28:195–200
- Couvillon PA, Bitterman ME (1992) A conventional conditioning analysis of "transitive inference" in pigeons. *J Exp Psychol Anim Behav Proc* 18:308–310
- D'Amato MR, Colombo M (1988) Representation of serial order in monkeys (*Cebus apella*). *J Exp Psych Anim Behav Proc* 14:131–139
- Davis H (1992) Transitive inference in rats (*Rattus norvegicus*). *J Comp Psychol* 106:342–349
- De Lillo C (1996) The serial organisation of behaviour by non-human primates: an evaluation of experimental paradigms. *Behav Brain Res* 81:1–17
- Delius J, Siemann M (1998) Transitive responding in animals and humans: exaptation rather than adaptation. *Behav Proc* 42:107–137
- Dusek JA, Eichenbaum H (1997). The hippocampus and memory for orderly stimulus relations. *Proc Natl Acad Sci USA* 94:7109–7114
- Fersen L von, Wynne CDL, Delius JD, Staddon JER (1991) Transitive inference formation in pigeons. *J Exp Psychol Anim Behav Proc* 17:334–341
- Gillan DJ (1981) Reasoning in chimpanzees II. Transitive inference. *J Exp Psychol Anim Behav Proc* 7:150–164
- Harris MR, McGonigle BO (1994) A model of transitive choice. *Q J Exp Psychol B* 47:319–348
- Higa JJ, Staddon ER (1993) "Transitive inference" in multiple conditional discriminations. *J Exp Anim Behav* 59:265–291
- Kallio KD (1982) Developmental change in a five-term transitive inference. *J Exp Child Psychol* 33:142–164
- Lazareva OF, Smirnova AA, Rayevsky VV, Zorina ZA (2000) Transitive inference in hooded crows: preliminary data. *Dokl Biol Sci* 370:30–32
- Luce RD (1959) Individual choice behavior. Wiley, New York
- McGonigle B, Chalmers M (1977) Are monkeys logical? *Nature* 267:694–696
- McGonigle B, Chalmers M (1992) Monkeys are rational! *Q J Exp Psychol B* 45:189–228
- Piaget J (1928) Judgement and reasoning in the child. Routledge, Chapman and Hall, London
- Rescorla RA, Wagner AR (1972) A theory of classical conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In: Black AH, Prokasy WF (eds) *Classical conditioning. II. Current research and theory*. Appleton-Century-Crofts, New York, pp 64–99
- Riley CA, Trabasso T (1974) Comparatives, logical structures and encoding in a transitive inference task. *J Exp Child Psychol* 17:187–203
- Roberts WA, Phelps MT (1994) Transitive inference in rats: a test of the spatial coding hypothesis. *Psychol Sci* 5:368–374
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. *Nature* 323:533–536
- Siemann M, Delius JD (1996) Influences of task concreteness upon transitive responding in humans. *Psychol Res* 59:81–93
- Siemann M, Delius JD (1998) Algebraic learning and neural network models for transitive and non-transitive responding. *Eur J Cogn Psychol* 10:307–334
- Steirn JN, Weaver JE, Zentall TR (1995) Transitive inference in pigeons: simplified procedures and a test of value transfer theory. *Anim Learn Behav* 23:76–82
- Terrace HS (1993) The phylogeny and ontogeny of serial memory: list learning by pigeons and monkeys. *Psychol Sci* 4:162–169
- Terrace HS, McGonigle B (1994) Memory and representation of serial order by children monkeys and pigeons. *Curr Dir Psychol Sci* 3:180–185
- Trabasso T (1977) The role of memory as a system in making transitive inferences. In: Kail RV, Hagen JW (eds) *Perspectives on the development of memory and cognition*. Lawrence Erlbaum, Hillsdale, pp 333–366
- Trabasso T, Riley C, Wilson E (1975) The representation of linear and spatial strategies in reasoning: a developmental study. In: Falmagne RJ (ed) *Reasoning: representation and process in children and adults*. Lawrence Erlbaum, Hillsdale, pp 201–229
- Treichler FR, Van Tilburg D (1999) Training requirements and retention characteristics of serial list organisation by macaque monkeys. *Anim Cogn* 2:235–244
- Weaver JE, Steirn JN, Zentall TR (1997) Transitive inference in pigeons: control for differential value transfer. *Psychonom Bull Rev* 4:113–117
- Wynne CDL (1995) Reinforcement accounts for transitive inference performance. *Anim Learn Behav* 23:207–217
- Wynne CDL (1997) Pigeon transitive inference: tests of simple accounts of a complex performance. *Behav Proc* 39:95–112
- Wynne CDL, Fersen L von, Staddon JER (1992) Pigeons' inferences are transitive and the outcome of elementary conditioning principles: a response. *J Exp Psychol Anim Behav Proc* 18:313–315